

Dell™ PowerVault™ Modular Disk Storage Manager User's Guide

Notes and Notices



NOTE: A NOTE indicates important information that helps you make better use of your computer.



NOTICE: A NOTICE indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

Information in this document is subject to change without notice.

© 2008 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge* and *PowerVault* are trademarks of Dell Inc.; *Microsoft*, *Windows*, *Windows Server*, *MS-DOS*, and *Internet Explorer* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *Red Hat* and *Red Hat Enterprise Linux* are registered trademarks of Red Hat, Inc. *SUSE* is a registered trademark of Novell Inc.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

February 2008

Contents

1	About This Guide	11
	User Interface	11
	Summary Tab	12
	Configure Tab	12
	Modify Tab	12
	Tools Tab	13
	iSCSI Tab	13
	Support Tab	13
	Other Information You May Need	14
2	About Your Storage Array	15
	Access Virtual Disk	15
	Out-of-Band and In-Band Management	16
	Adding Storage Arrays	17
	Automatic Discovery of Storage Arrays	17
	Manual Addition of a Storage Array	17
	Naming Storage Arrays	18
	Removing Storage Arrays	18
	Setting Up Your Storage Array	19
	Storage Array Support Data	21

Setting a Password	21
Password Guidelines	22
Resetting a Password	22
Connecting the Serial Cable	22
System Setup for Password Reset	23
Reset Password	24
Changing Expansion Enclosure ID Numbers	25
Configuring Alert Notifications	25
Configuring E-mail Alerts	25
Configuring SNMP Alerts	27
Starting or Restarting the Host-Agent Software in Windows	28
Starting or Restarting the Host-Agent Software in Linux	28
3 Using iSCSI	31
Using the iSCSI Tab	31
Changing the iSCSI Target Authentication	32
Entering Mutual Authentication Permissions	32
Creating CHAP Secrets	32
Changing the iSCSI Target Identification	34
Changing the iSCSI Target Discovery	34
Configuring the MD3000i iSCSI Host Ports	35
Advanced iSCSI Host Ports Settings	35
Viewing or Ending an iSCSI Session	36

	Viewing iSCSI Statistics and Setting Baseline Statistics	37
	Edit, Remove, or Rename Host Topology	38
4	Event Monitor	41
	Enabling the Event Monitor	41
	Disabling the Event Monitor	42
5	About Your Host	43
	Configuring Host Access	43
	Automatic Configuration	44
	Manual Configuration (using SAS HBA)	44
	Manual Configuration (using iSCSI)	45
	Removing Host Access	46
	Host Groups	47
	Creating a Host Group	47
	Adding a Host to a Host Group	47
	Removing a Host From a Host Group	48
	Moving a Host to a Different Host Group	48
	Removing a Host Group	49
	Host Topology	49
	Host Context Agent	50
	I/O Data Path Protection	51
	Failover with Red Hat Enterprise Linux	51

6	Disk Groups and Virtual Disks	53
	Creating Disk Groups and Virtual Disks	54
	Automatic Configuration	55
	Manual Configuration	55
	Hot Spare Drive Protection	58
	Automatically Configuring Hot Spares	58
	Manually Configuring Hot Spares	59
	Host-to-Virtual Disk Mapping	60
	Creating Host-to-Virtual Disk Mappings	60
	Modifying and Removing Host-to-Virtual Disk Mapping	60
	Changing Controller Ownership of the Virtual Disk	61
	Storage Partitioning	61
	Disk Group and Virtual Disk Expansion	62
	Disk Group Expansion	62
	Virtual Disk Expansion	62
	Storage Array Media Scan	63
	Changing Media Scan Settings	63
	Suspending the Media Scan	64
	Microsoft Services	64
	Virtual Disk Service	64
	Volume Shadow-Copy Service	65

7 Premium Feature—Snapshot
Virtual Disks 67

Creating a Snapshot Virtual Disk Using the Simple Path	68
About the Simple Path	68
Preparing Host Servers to Create the Snapshot Using the Simple Path	69
Creating the Snapshot Using the Simple Path	71
Creating a Snapshot Virtual Disk Using the Advanced Path	73
About the Advanced Path	73
Preparing Host Servers to Create the Snapshot Using the Advanced Path	74
Creating the Snapshot Using the Advanced Path	76
Specifying Snapshot Virtual Disk Names	78
Snapshot Repository Capacity	79
Re-creating Snapshot Virtual Disks	80
Disabling a Snapshot Virtual Disk	80
Preparing Host Servers to Re-create a Snapshot Virtual Disk	81
Re-creating a Snapshot Virtual Disk	82

8 Premium Feature—Virtual Disk
Copy 83

Creating a Virtual Disk Copy for an MSCS Shared Disk	84
Virtual Disk Read/Write Permissions	85
Virtual Disk Copy Restrictions	86

Creating a Virtual Disk Copy	87
Preparing Host Servers to Create a Virtual Disk Copy	87
Copying the Virtual Disk	88
Storage Array Performance During Virtual Disk Copy	89
Setting Copy Priority	89
Stopping a Virtual Disk Copy	90
Recopying a Virtual Disk	90
Preparing Host Servers to Recopy a Virtual Disk	91
Recopying the Virtual Disk	92
Removing Copy Pairs	93
9 Firmware Downloads	95
Downloading RAID Controller and NVSRAM Packages	95
Downloading Both RAID Controller and NVSRAM Firmware	96
Downloading Only NVSRAM Firmware	97
Downloading Non-redundant MSCS NVSRAM Firmware	97
Downloading Physical Disk Firmware	98
Downloading EMM Firmware	99
10 Troubleshooting Problems	101
Recovery Guru	101
Storage Array Profile	101

Device Health Conditions	101
SMrepassist Utility	102
Support Information Package	103
Unidentified Devices	104
Recovering from an Unidentified Storage Array	104
A Enclosure Hardware Replacement, Maintenance, and Configuration Considerations	107
Removing and Inserting Enclosure Management Modules on Attached Expansion Enclosures	107
Removing an EMM from the Expansion Enclosure	107
Inserting an EMM into an Expansion Enclosure	107
Removing and Inserting Physical Disks	108
MD3000 Maintenance Considerations	108
MD3000 Cluster Configuration Guidelines for Standalone Host Servers	109
Index	111

About This Guide

Dell™ PowerVault™ Modular Disk (MD) Storage Manager software is used to create and manage multiple storage arrays. The software can be used on any host attached to the storage array, as well as on storage management stations connected to the same sub-network.

MD Storage Manager is a graphical user interface (GUI) with wizard-guided tools and a task-based structure designed to reduce the complexity of installation, configuration, management, and diagnostic tasks.

MD Storage Manager software also contains an optional event monitoring service that is used to send alerts when a critical problem with the storage array occurs and a command line interface (CLI) to access functions performed by MD Storage Manager.

This guide is intended for users who are already familiar with the basic functions of their storage array. Any differences in certain functions between supported operating systems are explained where applicable.

MD Storage Manager online help contains detailed answers to software-related questions. You can access online help by clicking **Help** located at the top right corner of the MD Storage Manager interface. Refer to your storage array's *Installation Guide* for information on installing the MD Storage Manager.

User Interface

The Storage Manager screen is divided into three parts:

- The *Title Bar* at the top of the screen displays the name of the application and the Dell logo.
- Beneath the Title Bar is the *Array Selector*, listing the MD Storage Array that is currently selected. The icon next to the array's name indicates its condition. You can choose another array by clicking the down-arrow next to the array's name and highlighting a different array in the drop-down list. Links to the right of the array name let you add or remove arrays from the list of managed arrays. Links to the far right provide access to online help or close the Storage Manager.

- Beneath the Array Selector is the *Content Area*. Several tabs appear in this area to group the tasks you can perform on the selected array. When you click on a tab, the Content Area displays links for the tasks you can perform. The following sections list some of the tasks you can perform under each tab.

Summary Tab

- See the status of a storage array
- See the hardware components in a storage array
- See storage array capacity
- See hosts, mappings, and storage partitions
- See virtual disk groups and virtual disks
- Access links to online help, FAQs, and a tutorial about storage concepts

Configure Tab

- Configure host access
- Create a host group
- Create hot spares
- Create virtual disks
- Create snapshot virtual disks (if enabled)
- Create virtual disk copies (if enabled)
- Create host-to-virtual disk mappings

Modify Tab

- Modify the host topology
- Rename and delete virtual disks
- Add free capacity to a disk group
- Change virtual disk ownership and the preferred path of virtual disks
- Modify a snapshot virtual disk (if enabled)
- Manage virtual disk copies (if enabled)
- Edit host-to-virtual disk mappings

Tools Tab

- Rename a storage array
- Set or change a password
- View or enable premium features
- Turn on indicator lights
- Change enclosure ID numbers
- Set or change enclosure tags
- Set up e-mail alerts and SNMP alerts
- Synchronize controller clocks, change the network configuration, or reset the battery age
- Inherit system settings
- Change media scan settings

iSCSI Tab

- Set the authentication method supported by the target
- Define permissions for mutual authentication
- Set an alias for the target for identification
- Modify the discovery method for iSCSI targets
- Set the parameters for iSCSI host ports
- View or end iSCSI sessions
- View iSCSI statistics



NOTE: The iSCSI tab is shown only in the MD Storage Manager when the controllers contain iSCSI host ports.

Support Tab

- Recover from a failure
- Gather support information
- View the storage array profile
- Download RAID controller, NVSRAM, and physical disk firmware
- Manage RAID controllers

- View online help
- View the event log

Other Information You May Need



CAUTION: For complete regulatory and safety information, see your *Product Information Guide*. Warranty information may be included within this document or as a separate one.

- *Setting Up Your Dell PowerVault MD* provides an overview of setting up and cabling your storage array.
- *Dell™ PowerVault™ MD Installation Guide* provides installation and configuration instructions for both software and hardware.
- *Dell™ PowerVault™ MD Hardware Owner's Manual* provides information about the enclosure hardware.
- *Dell™ PowerVault™ MD Storage Manager CLI Guide* provides information about using the command line interface (CLI).
- *Dell™ PowerVault™ MD Resource CD* contains all system documentation and management tools.
- *Dell™ PowerVault™ MD Systems Support Matrix* provides information on supported software and hardware for MD systems. The document is available at support.dell.com.
- Dell PowerEdge® Cluster Documentation is available at support.dell.com. A link to clustering documentation is also included on the *Resource CD* under **Product Documentation**.
- Documentation for any components you purchased separately provides information to configure and install these options.
- Release notes or readme files are included to provide last-minute updates to the enclosure or documentation or advanced technical reference material intended for experienced users or technicians.



NOTE: Always check for updates on support.dell.com and read the updates first because they often supersede information in other documents.

About Your Storage Array


This chapter covers basic information about how to manage storage arrays from MD Storage Manager, including adding and removing arrays from the software, performing initial setup tasks, setting passwords on an array, and configuring alert notifications. For information on planning your storage array, see the *Installation Guide*.

Following is a list of terms that are used throughout this chapter:

- Physical Disk — Non-volatile, randomly-addressable device for storing data.
- Host — System that accesses a storage array and is mapped to virtual disks.
- Host Group — Hosts that are logically associated and share access to the same virtual disks.
- Host-Agent Software — Software installed on the host that provides in-band management and topology discovery.
- Logical Unit Number (LUN) — Address that identifies individual virtual disks within a storage array.
- Event Monitor — A feature that, when enabled, monitors activity on managed storage arrays and notifies a host or remote system when critical problems occur.
- SNMP Alert — Alert (SNMP trap) that is sent from the event monitor to an SNMP-enabled host.

Access Virtual Disk

Each RAID controller in an MD Storage Array maintains an *access virtual disk*. The host-agent software uses the access virtual disk to communicate management requests and event information between the storage management station and the RAID controller module in an in-band-managed storage array. The access virtual disk is not available for application data storage. The default LUN is 31.

-  **NOTICE:** Removing or manipulating an access virtual disk can cause a loss of management access. If you remove an access virtual disk mapping from an in-band-managed storage array, MD Storage Manager can no longer access the storage array. Do not modify the access virtual disk either in the operating system or with MD Storage Manager.


Out-of-Band and In-Band Management

You can manage a storage array in two ways:

- Out-of-band management
- In-band management


For out-of-band management, data is separate from commands and events. Data travels through the host-to-controller interface, while commands and events travel through the management port Ethernet cables.

When you use out-of-band management, you must set the network configuration for each RAID controller module's management Ethernet port, including its Internet Protocol (IP) address, subnetwork mask (subnet mask), and gateway. If you are using a Dynamic Host Configuration Protocol (DHCP) server, you can enable automatic network configuration, but if you are not using a DHCP server, you must enter the network configuration manually.

 **NOTE:** RAID controller module network configurations can be assigned using a DHCP server (the default setting). However, if a DHCP server is not available and the 10-second selection period times out, the RAID controller modules use the 192.168.128.101 static IP address for controller 0. For controller 1, the RAID controller modules use the 192.168.128.102 static IP address.

For in-band management, commands, events, and data travel through the host-to-controller interface. Unlike out-of-band management, commands and events are mixed with data.

For detailed information on setting up in-band and out-of-band management see the *Installation Guide*.

 **NOTE:** In-band management is not supported on systems running the Red Hat® Enterprise Linux® 3 operating system.

Adding Storage Arrays

To add a storage array to MD Storage Manager, click **New** in the Array Selector area. A window is displayed that allows you to choose the automatic or manual process to add a new storage array.



NOTE: Verify that your host or management station network configuration—including station IP address, subnet mask, and default gateway—is correct before adding a new storage array using the **Automatic** option.



NOTE: For Linux, set the default gateway so that broadcast packets are sent to 255.255.255.255. For Red Hat® Linux®, if no gateway exists on the network, set the default gateway to the IP address of the NIC.



NOTE: MD Storage Manager uses TCP/UDP port 2463 for communication to the MD Storage Array.

Automatic Discovery of Storage Arrays

The Automatic Discovery process sends out a broadcast message across the local subnetwork (subnet) and adds any storage array that responds to the message. The Automatic Discovery process finds both in-band and out-of-band storage arrays.

Manual Addition of a Storage Array

Use Manual Addition if the storage array resides outside of the local subnet. This process requires specific identification information to manually add a storage array:

- To add a storage array that uses in-band management, specify the host name or IP address of the host.

When adding a storage array using in-band management with iSCSI, a session must first be established between the initiator on the host server and the storage array. For more information, see "Configuring iSCSI" in the *Modular Disk 3000i Systems Installation Guide*.

The host agent must be restarted before in-band management communication can be established. See "Starting or Restarting the Host-Agent Software in Windows" on page 28 or "Starting or Restarting the Host-Agent Software in Linux" on page 28.

- To add a storage array that uses out-of-band management, specify the host name or IP address of each controller in the storage array.



NOTE: It can take several minutes for MD Storage Manager to connect to the specified storage array.

Naming Storage Arrays

Each storage array should be assigned a unique name. A storage array name has a 30-character limit. All leading and trailing spaces are deleted from the name. A name can consist of letters, numbers, and the special characters underscore (_), dash (-), and pound sign (#). No other special characters are allowed.

To physically locate a storage array:

- 1 Click the **Tools** tab.
- 2 Click **Blink** and then click **Blink Storage Array or Enclosures**.
- 3 Select the storage array from the displayed list and click **Blink**.
The indicator light on the front of the storage array flashes.
- 4 Click **Stop** after you locate the array.

To rename the selected storage array:

- 1 Click the **Tools** tab.
- 2 Click **Rename Storage Array**.
- 3 Type a unique, meaningful name that is easy to understand and remember.
- 4 Click **OK**.



NOTE: Avoid arbitrary names or names that might lose meaning in the future.

Removing Storage Arrays

You can remove a storage array from the list of managed arrays if you no longer want to manage it from a specific storage management station. Removing a storage array does not affect the storage array or its data in any way. Removing a storage array simply removes it from the list of storage arrays that appear in the drop-down list in the Array Selector. If a storage array is accidentally removed, it can be added again (see "Adding Storage Arrays" on page 17).

To remove a storage array:

- 1 Click **Remove** located to the right of the drop-down menu in the Array Selector.
- 2 Click **OK** in the **Remove** dialog box.

You can still manage the storage array from other storage management stations where it has been added.

Setting Up Your Storage Array

The **Perform Initial Setup Tasks** link located on the **Summary** tab provides links to the basic steps you should follow when initially setting up a storage array in MD Storage Manager. Following these steps ensures that you complete all the basic steps to configure your storage array.

Initial setup tasks include:

- 1 **Blink the Storage Array** — Find the physical location of the storage array on your network. The storage array can then be identified with a label.
- 2 **Rename the Storage Array** — Provide a unique and memorable name to help you easily identify the storage array.
- 3 **Set a Storage Array Password** — Set a unique password to prevent unapproved manipulation of the storage array, such as deletion of a virtual disk.
- 4 **Set up alert notifications** — Enable e-mail and SNMP alerts to notify administrators about storage array conditions that require attention. See "Configuring Alert Notifications" on page 25 for more information.
 - a **Configure Sender E-mail Settings** — Provide the SMTP, e-mail address, and contact information MD Storage Manager uses to send e-mail alerts.
 - b **Add or Edit E-mail Addresses** — Provide information about accounts that should receive e-mail-based alerts.
 - c **Set up SNMP Alerts** — Provide information about hosts that should receive SNMP-based alerts.
- 5 **Configure iSCSI Host Ports** — Configure network parameters for the iSCSI host ports on the RAID controller module(s).

- 6** Configure Host Access — Set up one or more hosts to access the storage array. See "Configuring Host Access" on page 43 for more information.
- 7** Configure storage array (2 options)
 - a** Automatic (Simple) configuration
 - Step 1: Automatic Configuration — See "Creating Disk Groups and Virtual Disks" on page 54 for more information.
 - Step 2: Create Host-to-Virtual Disk Mappings — See "Creating Host-to-Virtual Disk Mappings" on page 60 for more information.
 - b** Manual (Advanced) configuration
 - Step 1: Create Virtual Disks — See "Manual Configuration" on page 55 for more information.
 - Step 2: Configure Hot Spare Physical Disks — See "Manually Configuring Hot Spares" on page 59 for more information.
- 8** Manage iSCSI Settings — This option will be present only if your controllers contain iSCSI host ports.
 - a** Change Target Authentication — Choose the authentication methods and permissions (if required) for an initiator to access the target.
 - b** Enter Mutual Authentication Permissions — If initiators require mutual authentication, you can enter permissions for the target to access the initiator.
 - c** Change Target Identification — Define an alias for the target for easy identification.
 - d** Change Target Discovery — Configure parameters for how the target will be discovered on the network.
- 9** View and Enable Premium Features (Optional) — You may have purchased premium features, including snapshot virtual disks and virtual disk copies. See which premium features are currently available to you and enable these features if they are currently turned off.
- 10** Configure Ethernet Management Ports (Optional) — Configure network parameters for the Ethernet management ports managing a storage array for out-of-band Ethernet connections.

Storage Array Support Data

Aggregated support data can be generated for a storage array to aid in remote troubleshooting and issue analysis. To generate the support data report:

- 1 Click the **Support** tab, then click **Gather Support Information**.
- 2 Click **Browse** to display the **Collect All Support Data** dialog box.
- 3 In the **Save in** drop-down box, navigate to the location where you want the report saved.
- 4 Type a meaningful name in the **File name** text box and click **Save**.
- 5 Click **Start**.

A compressed (zip) file containing support data is saved to the location of your choice.

Setting a Password

You can configure each storage array with a password to protect it from unauthorized access. MD Storage Manager asks for this password when an attempt is made to change the storage array configuration, such as when a virtual disk is created or deleted. View operations do not require a password.

To set, change, or remove a password for a storage array:

- 1 Click the **Tools** tab, then click **Set or Change Password**.

Text boxes for the current password, the new password, and new password confirmation are displayed.

- 2 To enter a new password:
 - Leave the **Current password** text box blank.
 - Enter the new password in the **New password** and **Confirm new password** text boxes.

To change a password:

- Enter the current password in the **Current password** text box.
- Enter the new password in the **New password** and **Confirm new password** text boxes.

To remove a password:

- Enter the current password in the **Current password** text box.
- Leave the **New password** and **Confirm new password** text boxes blank.

If you forget your password, contact Dell for technical assistance.

Password Guidelines

Consider these guidelines when you create a password:

- Use secure passwords for your storage array. A password should be easy for you to remember but difficult for others to determine. Consider using numbers or special characters in the place of letters, such as a 1 in the place of the letter *l*, or the at sign (@) in the place of the letter *a*.
- For increased protection, use a long password with at least 15 alphanumeric characters. The maximum password length is 30 characters.
- Passwords are case sensitive.
- For security reasons, you can attempt to enter a password only ten times before the storage array enters a lockout state. Before you can try to enter a password again, you must wait ten minutes for the storage array to reset.

Resetting a Password

Perform this procedure when you have lost or forgotten your password and you need to reset it.

Connecting the Serial Cable

- 1 Remove the serial cable from the password reset cable package.
- 2 Connect the DB9 (oval) end of the cable to the serial port on the computer to be used to communicate with the RAID Controller module.
- 3 Connect the PS2-type (round) end of the cable to the serial port on either of the MD RAID Controller Modules. The flat side of the connector faces down when inserting.

System Setup for Password Reset

Microsoft® Windows® Operating Systems

- 1 Click Start → Programs → Accessories → Communication → HyperTerminal to run HyperTerminal.

If HyperTerminal is not installed, click Control Panel → Add/Remove Programs → Add/Remove Windows Components, find HyperTerminal and click the check-box, then click **Apply** and **OK**.



NOTE: The original Windows installation disk may be needed to install HyperTerminal.



NOTE: HyperTerminal is not a component on Windows Server® 2008 operating systems.

- 2 When HyperTerminal prompts for a name, type MD and click **OK**.
- 3 Select the COM1 port and click **OK**.
- 4 Set the following communication settings, click **Apply**, then click **OK**.

Bits per second: 115200

Data bits: 8

Parity: none

Stop bits: 1

Flow control: none

LINUX Operating System

The following instructions use the Linux application, MINICOM, to connect via the serial port:

- 1 Open a terminal/command window.
- 2 At the prompt, type `minicom` (all lowercase) and press <Enter>.
- 3 Once MINICOM is open, press <Ctrl><A>, then <Z>, then the letter <O> to open the configuration screen.
- 4 Select **Serial Port Setup** and press <Enter>.
- 5 Press <F> to change **Hardware Flow Control** setting to no.
- 6 Press <E> to set the **Comm Parameters**.

- 7 Press <I> to set the speed to 115200.
- 8 Press <Q> to set the data, parity, and stopbits to 8-N-1, then press <Enter>.
- 9 Press <Enter> to exit the **Comm Parameters** screen.
- 10 Select Exit and press <Enter> again to exit the setup screen.

Reset Password



NOTICE: Failure to stop data I/O to a non-fault-tolerant array before performing the following steps may result in loss of data.

- 1 Stop all I/O to the array.
- 2 From the HyperTerminal (Windows) or MINICOM (Linux) window, press <Ctrl>.
- 3 At the **SPECIAL OPERATIONS MENU**, press <1>, <0>, and <Enter>.
- 4 At the **SERIAL INTERFACE MODE MENU**, press <1> and <Enter>, then <Q> and <Enter>.
- 5 At the **BOOT OPERATIONS MENU**, press <R> and <Enter> to restart the controller. Text scrolls across the screen as the controller reboots.
- 6 Once Controller 0 has fully booted (look for `sodMain complete` in the HyperTerminal or MINICOM window), press <Enter>.
- 7 At the prompt, type `clearSYMBOLPassword` (must use exact case) and press <Enter>.
The return value = 0 = 0x0 indicates that the password has been reset /deleted.
- 8 Close HyperTerminal or MINICOM.
- 9 Remove password reset cable.
- 10 To set a new password, go to the Modular Disk Storage Manager software under the **Tools** tab and click **Set Or Change Password Link**.
- 11 Leave the Current Password blank, enter the new password twice, and click **OK**.



NOTE: If you require help with this procedure, contact Dell for technical assistance. For more information on contacting Dell, see the "Getting Help" chapter of the *Hardware Owners Manual*.

Changing Expansion Enclosure ID Numbers

When an MD1000 expansion enclosure is attached to an MD3000/MD3000i storage array for the first time, an enclosure ID number is assigned and maintained by the MD1000. This enclosure ID number is also shown in the MD Storage Manager, but it is not an indicator of the enclosure's physical location. It may appear that MD Storage Manager is reporting the expansion enclosures in improper order.

You can change the enclosure ID numbers in the MD Storage Manager by clicking the **Tools** menu and then clicking **Change Enclosures ID Numbers**. Any ID number you assign will not conflict with the enclosure IDs.

Configuring Alert Notifications

MD Storage Manager can send an alert for any condition on the storage array that requires your attention, such as the failure of a storage array component or the occurrence of an adverse environmental condition. Alerts can be sent as e-mail messages or as SNMP messages.

You can verify whether alerts are currently set by looking at the **Alert status** line in the **Status** area of the **Summary** tab.

Configuring E-mail Alerts

To configure e-mail alerts, click the **Tools** tab and then click **Set up e-mail alerts**. The **Content Area** displays two links: **Configure Sender E-mail Settings** and **Add or Edit E-mail Addresses**.



NOTE: These settings apply to all storage arrays currently managed by the management station.

Sender e-mail settings include the SMTP and e-mail address information MD Storage Manager uses to send e-mail alerts. To configure sender e-mail settings:

- 1 Click the **Tools** tab, then click **Set Up Email Alerts**.
- 2 Enter the following information:
 - **Sender email address** — The e-mail address that appears as the sender on every e-mail alert, such as that of the network administrator.
 - **Mail (SMTP) server** — The name of the Simple Mail Transfer Protocol (SMTP) gateway of the mail server from which e-mail alerts will be sent. For example, `smtp.mycompany.com`.
 - **Edit Sender Contact Information (Optional)** — Additional information about the sender such as the sender's name, company, and phone number. This information is optional; e-mail alerts will work if contact information is not provided.

To specify to whom alerts are sent:

- 1 Click the **Tools** tab, then click **Set Up Email Alerts**.
- 2 Type an e-mail address in the **Recipient email addresses** text box and click **Add** to add it to the list of configured e-mail addresses.
- 3 Specify the following for each e-mail address in the list (to choose a different setting, click the down arrow to the right of the field):

Information To Send — Select one of the following options from the drop-down list:

- **Event Only** — The alert e-mail contains only the event information. This alert type is the default.
- **Event + Profile** — The alert e-mail contains the event information and the storage array profile.
- **Event + Support** — The alert e-mail contains the event information and a compressed file that contains complete support information for the storage array that has generated the alert.

Frequency — Select one of the following options from the drop-down list:

- **Every event** — Sends an e-mail whenever an event occurs. This option is the default.

- **Every x hours** — Sends an e-mail at the specified interval if an event occurred during that period. You can select this option only if the **Information to send** option is set to **Event + Profile** or **Event + Support**.

4 Click **Save**.

Configuring SNMP Alerts

To add a management console to the list of addresses configured to receive SNMP alerts:



NOTE: The Management Information Base (MIB) for the MD Storage Array is copied to the client directory as part of a Full or Management Station installation selection. **DellMDStorageArray.mib** can be compiled on an SNMP Management Console using the interface provided by the console.

1 Click the **Tools** tab, then click **Set up SNMP Alerts**.

2 Enter the **Community name**.



NOTE: The community name is an ASCII string that identifies a known set of management consoles and is set by the network administrator in the management console. The default community name is public.

3 Enter the **Trap destination**.



NOTE: The trap destination is the IP address or the host name of a management console that runs an SNMP service.

4 Click **Add** to add the management console to the **Configured SNMP addresses** list.

5 Repeat steps 2 through 4 until you have added all management consoles that should receive SNMP alerts.

6 Click **OK**.



NOTE: You must install an SNMP service on every system included in the list of addresses configured to receive SNMP alerts.



NOTE: You do not have to install MD Storage Manager on a system in order to display SNMP alerts. You need only install an appropriate SNMP service and application (such as the Dell IT Assistant).

Starting or Restarting the Host-Agent Software in Windows

The SMagent software automatically starts after you reboot the host. If you add a storage array after the host server has started, or if iSCSI sessions are created while the SMagent is running, you must restart the SMagent software manually using the following procedure:

- 1 Click **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**.
or
Click **Start** → **Administrative Tools** → **Services**.
- 2 In the Services dialog, select **Modular Disk Storage Manager Agent**.
- 3 If the Modular Disk Storage Manager Agent is running, click **Action** → **Stop**, then wait approximately 5 seconds.
- 4 Click **Action** → **Start**.

Starting or Restarting the Host-Agent Software in Linux

The SMagent software automatically starts after you reboot the host. If you add a storage array after the host server has started, or if iSCSI sessions are created while the SMagent is running, you must restart the SMagent software manually using the following procedure.

To start or restart the host-agent software, enter the following command at the prompt:

```
SMagent start
```

The SMagent software might take a little time to initialize. The cursor is shown, but the terminal window does not respond. When the program starts, the following message is displayed:

```
SMagent started.
```

After the program completes the startup process, text similar to the following messages is displayed:

```
Storage Manager Agent, Version 09.1x.00.00 Built  
Wed Aug 15
```

```
16:54:46 CDT 2006 Copyright (C) 2006. All rights  
reserved.
```

```
checking device /dev/rdisk/c0t0d0s2 : skipping  
checking device
```

```
/dev/rdisk/c2t3d18s2 : skipping checking device
```

```
/dev/rdisk/c2t3e16s2 : skipping checking device
```

```
/dev/rdisk/c2t3d14w2 : skipping
```


Using iSCSI

This chapter provides information on using iSCSI in MD Storage Manager. For iSCSI prerequisite requirements and detailed step-by-step instructions on setting up and configuring iSCSI, see the *Installation Guide*.


Using the iSCSI Tab

The iSCSI tab is shown in the MD Storage Manager only when the controllers contain iSCSI host ports. You can define or change settings for the iSCSI target or enter the CHAP permissions in the iSCSI tab. Here are some of the iSCSI settings:

- **Change Target Authentication** — Select the authentication method to be supported by the target.
- **Enter Mutual Authentication Permissions** — Define the permissions for initiators that require mutual authentication.
- **Change Target Identification** — Associate an alias with the target for simpler identification.
- **Change Target Discovery** — Modify the way to discover iSCSI targets using the Internet Storage Name Service (iSNS) server settings.
- **Configure iSCSI Host Ports** — Set the parameters for iSCSI host ports.
- **View/End iSCSI Sessions** — View iSCSI session details and end iSCSI sessions.
- **View iSCSI Statistics** — View and save iSCSI statistics.

Changing the iSCSI Target Authentication


If an initiator requires mutual (bi-directional) authentication see "Entering Mutual Authentication Permissions" on page 32.

- 1 Click the **iSCSI** tab, and then click **Change Target Authentication**.
- 2 Select **None** if no authentication is required for any initiator to access the target.
 **NOTE:** If you select **None**, any initiator can access this target. Use this option only if you do not require secure data. However, if you select both **None** and **CHAP** at the same time, the storage array will allow an iSCSI initiator to log on with or without CHAP authentication.
- 3 Select **CHAP** if you want any initiator that tries to access the target to provide the target permissions. If **CHAP** is selected, but no CHAP secret is defined, an error message appears.

Click **CHAP Secret** to see the **Enter CHAP Secret** dialog (see "Creating CHAP Secrets" on page 32). You can define the permissions in this dialog.

Entering Mutual Authentication Permissions

Mutual authentication or two-way authentication is a way for a client or a user to verify themselves to a host server, and for the host server to validate itself to the user. This validation is accomplished in such a way that both parties are sure of the other's identity.

- 1 Click the **iSCSI** tab, and then click **Enter Mutual Authentication Permissions**.
- 2 Select an initiator from the list. The initiator details are shown.
- 3 Select **CHAP Secret** to enter the initiator CHAP permissions in the dialog that appears.
 **NOTE:** To add, modify, or delete an initiator, click the **Modify** tab, and then click **Edit Host Topology**.

Creating CHAP Secrets

When you set up an authentication method, you can choose to create a Challenge Handshake Authentication Protocol (CHAP) secret. The CHAP secret is a password that is recognized by the initiator and the target. If you are using mutual authentication to configure the MD3000i storage array, you

must enter the same CHAP secret that is defined in the iSCSI initiator, and you must define a CHAP secret on the target (the storage array) that must be configured in every iSCSI initiator that will connect to the target. For more information on CHAP, see "Understanding CHAP Authentication" in the *Installation Guide*.

Initiator CHAP Secret

The initiator CHAP secret is set on the host using the iSCSI initiator. If you are using the mutual authentication method, you must define the initiator CHAP secret when you set up the host. This must be the same CHAP secret that is defined for the target when defining mutual authentication settings.

Target CHAP Secret

If you are using CHAP secrets, you must define the CHAP secret for the target.

Valid Characters for CHAP Secrets



The CHAP secret must be between 12 and 57 characters. The CHAP secret supports characters with ASCII values of 32 to 126 decimal. See Table 3-1 for a list of valid ASCII characters.

Table 3-1. Valid ASCII Characters for CHAP Secrets

Space	!	"	#	\$	%	&	'	()	*	+
,	-	.	/	0	1	2	3	4	5	6	7
8	9	:	;	<	=	>	?	@	A	B	C
D	E	F	G	H	I	J	K	L	M	N	O
P	Q	R	S	T	U	V	W	X	Y	Z	[
\]	^	_	a	b	c	d	e	f	g	h
i	j	k	l	m	n	o	p	q	r	s	t
u	v	w	x	y	z	{		}	~		

Changing the iSCSI Target Identification

You cannot change the iSCSI target name, but you can associate an alias with the target for simpler identification. Aliases are useful because the iSCSI target names are not intuitive. You should provide an iSCSI target alias that is meaningful and easy to remember.

- 1 Click the **iSCSI** tab, and then click **Change Target Identification**.
- 2 Type the alias in the **iSCSI target alias** field and click **OK**.
 -  **NOTE:** Aliases can contain a maximum of 30 characters. Aliases can include letters, numbers, and the special characters underscore (`_`), minus (`-`), and pound sign (`#`). No other special characters are permitted.
 -  **NOTE:** Open iSCSI (which is used by Red Hat Enterprise Linux 5 and SUSE® Linux Enterprise Server 10 with SP 1) does not support using target alias.

Changing the iSCSI Target Discovery

- 1 Click the **iSCSI** tab, and then click **Change Target Discovery**.
- 2 Select the **Use iSNS server** check box to activate iSCSI target discovery. You can use *one* of these methods:
 - a Use the **DHCP option (IPv4 only)** to automatically activate target discovery. You also can refresh the **DHCP**.
 - b Type the **IPv4 or IPv6 address** to activate the target discovery. After you manually enter an IP address, you also can click **Advanced** to set the customized **TCP listening ports**.
- 3 If you do not want to allow discovery sessions that are not named, select **Disallow un-named discovery sessions**.

Un-named discovery sessions are discovery sessions that are permitted to run without a target name. With an un-named discovery session, the target name or the target portal group tag is not available to enforce the iSCSI session identifier (ISID) rule. For more information on un-named discovery sessions, click the **Support** tab, then click **View Online Help**.

Configuring the MD3000i iSCSI Host Ports

Use the configuration dialog for the iSCSI host ports to set up the MD3000i iSCSI host ports to use with storage arrays in a storage area network (SAN).

- 1 Click the **iSCSI** tab, and then click **Configure iSCSI Host Ports**.
- 2 Select the controller in the **iSCSI host port** field, and then use *one* of these methods to configure the port:
 - a Automatically obtain the configuration using one of the following methods:
 - IPv4 — Obtain the configuration from the DHCP server, or refresh DHCP.
 - IPv6 — Obtain the configuration automatically from a router.
 - b Manually specify the configuration using one of the following methods:
 - IPv4 — Enter the IP address, subnet mask, and gateway for the host port.
 - IPv6 — Enter the IP address, routable IP addresses, and router IP address.

After you manually enter an IP address, you also can click **Advanced** to set the advanced parameters for the iSCSI target discovery.

Advanced iSCSI Host Ports Settings



NOTE: Configuring the advanced iSCSI host ports settings is optional.

Use the advanced settings for the individual iSCSI host ports to specify the TCP frame size, the virtual LAN, and the network priority.

Table 3-2. Advanced iSCSI Host Port Settings

Setting	Description
Virtual LAN (VLAN)	<p>A method of creating independent logical networks within a physical network. Several VLANs can exist within a network. VLAN 1 is the default VLAN.</p> <p>NOTE: For more information on creating and configuring a VLAN with MD Support Manager, click the Support tab, then click View Online Help.</p>

Table 3-2. Advanced iSCSI Host Port Settings (continued)

Setting	Description
Ethernet Priority	The network priority can be set from lowest to highest. Although network managers must determine these mappings, the IEEE has made broad recommendations: <ul style="list-style-type: none">• 0 — lowest priority (default)• 1-4 — ranges from "loss eligible" traffic to controlled-load applications, such as streaming multimedia and business-critical traffic• 5-6 — delay-sensitive applications such as interactive video and voice• 7 — highest priority reserved for network-critical traffic (do not use with the MD3000i)
TCP Listening Port	The default Transmission Control Protocol (TCP) listening port is 3260.
Jumbo Frames	The maximum transmission units (MTUs). It can be set between 1500 and 9000 bytes per frame. If the Jumbo Frames are disabled, the default MTU is 1500 bytes per frame.





NOTE: Changing any of these settings resets the iSCSI port. I/O is interrupted to any host accessing that port. You can access the I/O automatically after the port restarts and the host logs in again.

Viewing or Ending an iSCSI Session

- 1 Click the iSCSI tab, and then click **View/End iSCSI Sessions**.
- 2 Select the session you want to view in the **Current sessions** box. The details are shown below in the **Details** box.
- 3 If you want to end the session, perform the following steps:
 - a Select the session that you want to end, and then click **End Session** to show the **End Session** confirmation window.

- b** In the confirmation window, type **yes** to confirm that you want to end the iSCSI session, and then click **OK**.

 **NOTE:** If you end a session, any corresponding connections terminate the link between the host and the storage array, and the data on the storage array is no longer available.

 **NOTE:** When a session is manually terminated using the MD Storage Manager, the iSCSI initiator software will automatically attempt to re-establish the terminated connection to the storage array. This may cause an error message.

- 4** Click **Save As** to save the entire iSCSI sessions topology as a text file.

Viewing iSCSI Statistics and Setting Baseline Statistics

If the configured storage array has iSCSI technology, the **View iSCSI Statistics** option is available only on the **iSCSI** tab.

- 1** Click the **iSCSI** tab, and then click **View iSCSI Statistics**.
- 2** Select the iSCSI statistic type you want to view. Select one of these types:
 - Ethernet MAC statistics
 - Ethernet TCP/IP statistics
 - Target (protocol) statistics

- 3** Choose either **Raw statistics** or the **Baseline statistics**.

Raw statistics are all the statistics that have been gathered since the controllers were started. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

After you select the statistics type and either raw or baseline statistics, the details of the statistics appear in the statistics tables.

- 4** To set the baseline for the statistics, complete the following steps:
 - a** Select **Baseline Statistics**.
 - b** Click **Set Baseline**.

- c Confirm that you want to set the baseline statistics in the dialog that appears.

The baseline time shows the latest time you set the baseline. The sampling interval is the difference in time from when you set the baseline until you launch the dialog or click **Refresh**.



NOTE: You must first set a baseline before you can compare baseline statistics.

Edit, Remove, or Rename Host Topology

If you give access to the wrong host or the wrong host group, you can remove or edit the host topology. Use one of the following actions to correct the host topology:

Table 3-3. Host Topology Actions

Desired Action	Steps to Complete Action
Move the host, the host group, or the iSCSI initiator.	<ol style="list-style-type: none"> 1 Click the Modify tab, and then click Edit Host Topology. 2 Select the item that you want to move, and then click Move. 3 Select a host group to move the host to and click OK.
Manually change the host type.	<ol style="list-style-type: none"> 1 Click the Modify tab, and then click Edit Host Topology. 2 Select the host that you want to change and click Change. 3 Select a new host type (or operating system) and click OK.
Manually delete the host, the host group, or the iSCSI initiator.	<ol style="list-style-type: none"> 1 Click the Modify tab, and then click Edit Host Topology. 2 Select the item that you want to remove and click Remove.
Rename the host, the host group, or the iSCSI initiator.	<ol style="list-style-type: none"> 1 Click the Modify tab, and then click Edit Host Topology. 2 Select the item that you want to rename and click Rename. 3 Type a new label for the host and click OK.

Table 3-3. Host Topology Actions (continued)

Desired Action	Steps to Complete Action
Add an iSCSI Initiator.	<ol style="list-style-type: none">1 Click the Modify tab, and then click Edit Host Topology.2 Select the host you want to add an iSCSI initiator to and click Add.3 Type or select an iSCSI initiator name and label for the iSCSI initiator and click OK.

Event Monitor

An event monitor is provided with MD Storage Manager. When enabled, the event monitor runs continuously in the background and monitors activity on the managed storage arrays. If the event monitor detects any critical problems, it can notify a host or remote system using e-mail, Simple Network Management Protocol (SNMP) trap messages, or both.

For the most timely and continuous notification of events, enable the event monitor on a management station that runs 24 hours a day. Enabling the event monitor on multiple systems or having a combination of an event monitor and MD Storage Manager active can result in duplicate events, but this does not indicate multiple failures on the array.

Enabling the Event Monitor

You can enable the event monitor at any time.



NOTE: It is a good idea to configure the event monitor to start by default on a management station that runs 24 hours a day.

Microsoft® Windows®

- 1 Select **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**.

or

Select **Start** → **Administrative Tools** → **Services**.

- 2 From the list of services, select **Modular Disk Storage Manager Event Monitor**.
- 3 Select **Action** → **Properties**.
- 4 In the **Service Status** area, click **Start**.

Linux

At the command prompt, type `SMmonitor start` and press <Enter>. When the program startup begins, the system displays the following message:

```
SMmonitor started.
```

Disabling the Event Monitor

Disable the event monitor if you do not want the system to send alert notifications. If you are running the event monitor on multiple systems, disabling the event monitor on all but one system prevents the sending of duplicate messages.

Windows

- 1 Select **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**.
or
Select **Start** → **Administrative Tools** → **Services**.
- 2 From the list of services, select **Modular Disk Storage Manager Event Monitor**.
- 3 Select **Action** → **Properties**.
- 4 In the **Service Status** area, click **Stop**.

Linux

At the command prompt, type `SMmonitor stop` and press <Enter>. When the program shutdown is complete, the system displays the following message:

```
Stopping Monitor process.
```

About Your Host

This chapter covers basic information about configuring host groups and host access, host topology, and I/O data path protection.

A host is a system that accesses a storage array and is mapped to the virtual disks through one or more host connections. Hosts have the following attributes:

- Host name — A name that uniquely identifies the host.
- Host type — The operating system running on the host.
- Host connection — A physical connection to the host server. Host connections can be automatically detected by MD Storage Manager and can be identified by an alias assigned by the user.
- Host group — A host may be associated with other hosts to share access to the same virtual disks.

Configuring Host Access

Configuring host access allows you to either permit or deny access to a storage array for specific hosts. When you permit host access, that host can then be mapped to a virtual disk on the storage array. On the **Summary** tab, the **Hosts & Mappings** area indicates how many hosts are configured to access the array. Click **Configured Hosts** in this area to see the names of these hosts.

Host access configuration is the first step in setting up your storage array. You must complete this task during initial setup and anytime you connect a new host.

After you configure host access, the host does not yet have the ability to write data to the storage array. You must map hosts to the virtual disks and register virtual disks with the host's operating system before a host can write to the storage array. See “Disk Groups and Virtual Disks” on page 53 for information on these tasks.

To begin configuring host access, click the **Configure** tab and then click either **Configure Host Access (Automatic)** or **Configure Host Access (Manual)**. See the appropriate section for manual configuration, depending on whether you are using SAS HBA or iSCSI.

Automatic Configuration

To automatically configure a host for access to the storage array:

- 1 Click the **Configure** tab and then click **Configure Host Access (Automatic)**.
- 2 To see hosts that already have access to the storage array, click **View configured hosts**.
- 3 Select the hosts you want to give access to the storage array in the **Available hosts** window.
- 4 To see the ports and the host type for the selected hosts, click **View Details** at the right of the list.
- 5 Click **Add** to move specific hosts to the **Selected hosts** window.
- 6 Click **OK** to configure access for the hosts you selected.

Manual Configuration (using SAS HBA)



NOTE: Host access that is manually configured requires special attention to ensure that the correct SAS host port World Wide IDs are selected for each host. If any incorrect IDs are configured, an inaccurate topology will result. You can use the SAS/5/E HBA BIOS Setup program to identify the World Wide IDs for the SAS host ports.

Configure the host to make it available to the storage array for volume mapping by following these steps.

- 1 Click the **Configure** tab and then click **Configure Host Access (Manual)**.
- 2 Type a name of your choice in the **Enter host name** text box.
This can be an informal name, not necessarily a name used to identify the host to the network.
- 3 Select the operating system of your host in the **Select host type** box and then click **Next**.

- 4 Specify the HBA host ports by choosing known host ports or by manually defining host ports.

To select a host port that is already recognized by MD Storage Manager, click a host port in the **Known HBA host ports** list, then click **Add**.

To manually define a host port, click **New**, enter the **HBA host port** and **Alias** in the **Enter New HBA Host Port** dialog box, and then click **Add**.

- 5 Click **Next**.

- 6 Indicate whether the host is part of a host group (cluster):

If the host is not part of a host group, select **No**.

If the host is part of a host group, select **Yes**:

- To create a new host group, enter a name in the **Enter new host group name** text box.
- To add the host to an existing host group, select the host group from the **Select existing host group** box.

- 7 Click **Next**.

- 8 Click **Finish** to configure the host.

Manual Configuration (using iSCSI)

Configure the host to make it available to the storage array for volume mapping by following these steps.

- 1 Click the **Configure** tab and then click **Configure Host Access (Manual)**.

- 2 Type a name of your choice in the **Enter host name** text box.

This can be an informal name, not necessarily a name used to identify the host to the network.

- 3 Select the operating system of your host in the **Select host type** drop-down box and then click **Next**.

- 4 Specify the iSCSI initiators by choosing known initiators or by manually defining initiators.

To select an initiator that is already recognized by MD Storage Manager, click an initiator in the **Known iSCSI Initiators** list, and then click **Add**.

To manually define an initiator, click **New**, enter the **iSCSI initiator name** and **iSCSI initiator label** in the **Enter new iSCSI initiator** dialog box, and then click **Add**.



NOTE: The initiator name entered must match the name on a host server that will connect to the storage array.

- 5 Click **Next**.
- 6 Indicate whether the host is part of a host group (cluster):
If the host is not part of a host group, select **No**.
If the host is part of a host group, select **Yes**:
 - To create a new host group, enter a name in the **Enter new host group name** text box.
 - To add the host to an existing host group, select the host group from the **Select existing host group** box.
- 7 Click **Next**.
- 8 Click **Finish** to configure the host.

Removing Host Access

Use the following procedure to remove a host's access to a storage array:

- 1 Click the **Modify** tab, then click **Edit topology**.
- 2 In the host topology list, click the plus sign (+) to the left of the host group name.
The host group expands to show the hosts in the group.
- 3 In the list, click the name of the host whose access you want to remove, and then click **Remove** located to the right of the list. Click **Yes** to remove access.
- 4 Repeat step 3 for each host whose access you want to remove.
- 5 When the list contains only those hosts you want to access the storage array, click **Close** beneath the list.

Host Groups

A host group is a logical entity of two or more hosts that share access to specific virtual disks on the storage array. You create host groups with MD Storage Manager.

All hosts in a host group must have the same host type (operating system). In addition, all hosts in the host group must have special software, such as clustering software, to manage virtual disk sharing and accessibility.

If a host is part of a cluster, every host in the cluster must be connected to the storage array, and every host in the cluster must be added to the host group.

Use the following procedures to create a host group, to add or remove hosts from a host group, or to delete a host group.

Creating a Host Group

- 1 Click the **Configure** tab and then click **Create Host Group**.
- 2 Type a name for the new host group in the **Enter new host group name** text box.
- 3 In the **Select hosts to add** list, click the name of a host you want to add to the host group, then click **Add**.
The host moves to the **Hosts in group** list.
- 4 Repeat step 3 until all the hosts you want to add to the host group are moved into the **Hosts in group** list.
- 5 Click **OK**.

Adding a Host to a Host Group

- 1 Click the **Modify** tab, then click **Edit Host Topology**.
A list of hosts and host groups appears.
- 2 In the host topology list, click the plus sign (+) to the left of the host group name.
The host group expands to show the hosts in the group.
- 3 Click the host you want to move and click **Move**.

- 4 Select the host group to which you want to move the host.
- 5 Click **OK**.

The host is moved into the host group.

The host retains the virtual disk mappings assigned to it, and inherits the virtual disk mappings assigned to the group. Other hosts in the group do not inherit the mappings of the added host.

Removing a Host From a Host Group

- 1 Click the **Modify** tab, then click **Edit Host Topology**.
A list of hosts and host groups appears.
- 2 In the host topology list, click the plus sign (+) to the left of the host group name.
The host group expands to show the hosts in the group.
- 3 Click the name of the host you want to remove from the group.
- 4 Click **Remove** located to the right of the list.
- 5 Click **Yes** to remove the host.

The host is moved out of the host group. The host retains the virtual disk mappings assigned to it, and loses the virtual disk mappings assigned to the group.

Moving a Host to a Different Host Group

- 1 Click the **Modify** tab, then click **Edit Host Topology**.
A list of hosts and host groups appears.
- 2 In the host topology list, click the plus sign (+) to the left of the host group name.
The host group expands to show the hosts in the group.
- 3 Click the name of the host you want to move to another group and click **Move**.
- 4 Select the host group to which you want to move the host.
- 5 Click **OK**.
The host is moved to the indicated host group.

The host retains the virtual disk mappings assigned to it, and inherits the virtual disk mappings assigned to the group to which it is moved. The host loses the virtual disk mappings assigned to the group from which it was moved.

Removing a Host Group

This section covers removing an entire host group. To remove a single host from a host group, see "Removing a Host From a Host Group" on page 48.

- 1 Click the **Modify** tab, then click **Edit Host Topology**.
- 2 In the host topology list, click the name of the host group you want to remove.
- 3 Click **Remove**.
- 4 Click **Yes**.

The host group and its assigned virtual disk mappings are removed.



NOTE: If the host group contains hosts, those hosts are removed as well, including their access to the storage array.

Host Topology

Host topology is the organization of hosts, host groups, and host interfaces configured for a storage array. The **Edit Host Topology** screen accessed from the **Modify** tab shows the hierarchy of the host groups, the hosts that are part of each host group, and the host connections of each host.

You can use these tasks to change the host topology:

- Move a host or a host connection
- Rename a host group, a host, or a host connection
- Add a host connection
- Replace a host connection
- Change a host type

MD Storage Manager automatically detects these changes for any host running the host agent software.

Host Context Agent

The host context agent discovers the host topology. The host context agent starts when the host is started and stops when the host is turned off. The topology discovered by the host context agent can be viewed by clicking **Configure Host Access (Automatic)** in the **Configure** tab in the MD Storage Manager.

You must stop and restart the host context agent to see the changes to the host topology if any of the following situations occur:

- A new storage array is attached to the host server.
- A host is added while turning on power to the RAID controller modules.

Linux

In Linux, you can stop and start the host context agent from the command line. Use the following syntax: `SMagent start` or `SMagent stop`.

You will stop and then restart `SMagent` after performing either of the two following maintenance tasks.

- Moving a controller offline or replacing a controller.
- Removing host-to-array connections from or attaching host-to-array connections to a Linux host server.

Windows

In Windows, you can stop and start the host context agent from the **Services** option of the **Administrative Tools**. To access the host context agent:

- 1 Select **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**.

or

Select **Start** → **Administrative Tools** → **Services**.

- 2 From the list of services, select **Modular Disk Storage Manager Agent**.

I/O Data Path Protection

You can have multiple host-to-array connections for a host. Make sure to select all of the connections to the array when configuring host access to the storage array.



NOTICE: Refer to the *Installation Guide* for more information on cabling configurations.



NOTE: For maximum redundancy, you must select all host connections to the array when manually defining host topology. For example, a host might have two host connections. For this host, you would select two host connections.

If a component such as a RAID controller module or a cable fails, or an error occurs on the data path to the preferred RAID controller module, virtual disk ownership is moved to the alternate nonpreferred RAID controller module for processing. This failure or error is called *failover*.

Multi-path drivers such as MPIO and MPP are installed on host systems that access the storage array and provide I/O path failover. The multi-path driver (MPIO in Windows and MPP in Linux) is used for failover. Automatic Virtual Disk Transfer (AVT) is used specifically for single-port cluster failover. The AVT feature mode is automatically selected by host type.



NOTE: You should have the multi-path driver installed on the hosts at all times, even in a configuration where there is only one path to the storage system, such as a single port cluster configuration.

During a failover, the virtual disk transfer is logged as a critical event, and an alert notification is sent automatically if you have configured alert destinations for the storage array.

Failover with Red Hat Enterprise Linux

For users running Red Hat Enterprise Linux Version 3.0 (x86), the `mpp_vhba` failover driver component does not load unless a storage array is connected and a LUN is mapped to the host server. Follow these steps to ensure proper MPP driver loading and LUN mapping:

- 1 Ensure that the storage array is connected to the host server and that out-of-band (Ethernet) connections to each RAID controller module exist.
- 2 Create a virtual disk and map it to LUN 0.
- 3 Reboot the host.

After the host is rebooted, the LUN is properly mapped and the MPP driver is loaded. All other sequential LUNs can be created and registered to the operating system by running **hot_add**.


Disk Groups and Virtual Disks


Following is a list of terms used throughout this chapter:

- **Disk Group** — A set of physical disks that are logically grouped and assigned a RAID level. Every disk group provides the overall capacity required to create one or more virtual disks.
- **Virtual Disk** — A logical component created to enable hosts to access storage on the storage array. A virtual disk is created from the capacity available on a disk group and appears as one logical component even though it is created from more than one physical disk.
- **Storage Partitioning** — Logical division of a storage array into entities consisting of one or more virtual disks that can be accessed by a single host or shared among hosts that are part of a host group.
- **Unconfigured Capacity** — Physical disks that are not already assigned to a disk group.
- **Free Capacity** — Space in a disk group that has not been assigned to a virtual disk.
- **Standby Hot Spare Drive** — Physical disk that has been assigned as a hot spare drive and is available to take over for any failed physical disk.
- **In-use Hot Spare Drive** — Physical disk that has been assigned as a hot spare drive and is currently taking over for a failed physical disk.
- **Snapshot Virtual Disk** — Point-in-time image of a virtual disk in a storage array.
- **Snapshot Repository Virtual Disk** — Virtual disk containing metadata and copy-on-write data for a particular snapshot virtual disk; automatically created when the snapshot virtual disk is created.
- **Consistency Check** — Background operation that checks the parity of virtual disks.

Creating Disk Groups and Virtual Disks

Disk groups are created in the unconfigured capacity of a storage array, and virtual disks are created in the free capacity of a disk group. The maximum number of physical disks supported in a disk group is 30. The hosts attached to the storage array read and write data to the virtual disks.

 **NOTE:** Before you can create virtual disks, you must first organize the physical disks into disk groups and configure host access. Then you can create virtual disks within a disk group.

 **NOTE:** The disk group must contain physical disks of the same type. Mixing SAS and SATA II disks in a disk group is not supported.

To create a virtual disk, use one of the following methods:

- Create a new disk group from unconfigured capacity. You first define the RAID level and free capacity (available storage space) for the disk group, and then you define the parameters for the first virtual disk in the new disk group.
- Create a new virtual disk in the free capacity of an existing disk group. You only need to specify the parameters for the new virtual disk.

A disk group has a set amount of free capacity that was configured when the disk group was created. You can use that free capacity to subdivide the disk group into one or more virtual disks.

You can create disk groups and virtual disks using an automatic configuration procedure or using a manual configuration procedure. Automatic configuration provides the fastest method, but with limited configuration options. Manual configuration is a more involved process, but provides more configuration options.

When creating a virtual disk, consider all of the possible uses for that virtual disk, and select an appropriate capacity for those uses. For example, if a disk group has a virtual disk that stores multimedia files (which tend to be large) and another virtual disk that stores text files (which tend to be small), the multimedia file virtual disk requires more capacity than the text file virtual disk.

A disk group should be organized according to its related tasks and subtasks. For example, if you create a disk group for the Accounting Department, you can create virtual disks that match the different types of accounting

performed in the department: Accounts Receivable (AR), Accounts Payable (AP), internal billing, and so forth. In this scenario, the AR and AP virtual disks probably need more capacity than the internal billing virtual disk.



NOTE: In Linux, the host must be rebooted after deleting virtual disks to reset the /dev entries.



NOTE: Before you can use a virtual disk, you must register the disk with the host systems. This process is described in “Host-to-Virtual Disk Mapping” on page 60.

Automatic Configuration

If you want to set up virtual disks quickly, click the **Configure** tab and then click **Automatic Configuration**. With this option, you do not need to configure individual options for each virtual disk. When you use automatic configuration:

- All available unconfigured capacity on the array is used.
- All disk groups have the same RAID level.
- All virtual disks have equal capacity.
- The number of virtual disks created is based on the selected RAID level and available unconfigured capacity.
- For a RAID level 1 or 5 disk group, hot spare drives are selected automatically based on the number of drives and types of drives available in the storage array.

Manual Configuration

To create individual virtual disks or disk groups, click the **Configure** tab and then click **Create Virtual Disks**. You create one disk group and virtual disk at a time, but have control over the RAID level and capacity for each virtual disk and disk group. Use this method if you have unique capacity requirements for a disk group or virtual disk.

Creating a Disk Group and Virtual Disk From Unconfigured Capacity

Use the following procedure to manually create a disk group:

- 1 Click the **Configure** tab, then click **Create Virtual Disks**.
- 2 Select **Unconfigured capacity** on the **Create Virtual Disks – Select Capacity Type** page, then click **Next**.

- 3 Select **Manual** on the **Create Virtual Disks – Physical Disk Selection Choices** page, then click **Next**.
- 4 Select the RAID level for the new disk group.
- 5 To select one physical disk to add to the disk group, click the disk of your choice in the **Unselected physical disks** list.

To select more than one physical disk to add to the disk group, press **<Ctrl>** while clicking the disks of your choice in the **Unselected physical disks** list.



NOTE: When adding more than one physical disk to a disk group, it is recommended to use disks with the same capacity. You can choose to use disks of differing capacities; however, the overall capacity of the disk group will be based on the smallest capacity physical disk. This means that additional capacity on larger physical disks will not be available for use.

- 6 Click **Add** to add the disk(s) you selected in step 5 to the **Selected physical disks** list.
- 7 Click **Calculate Capacity** beneath the list of selected disks to see the capacity of the disk group you are creating.
- 8 To add or remove capacity for the proposed disk group, highlight disks in either list and click **Add** or **Remove**.
- 9 When you are satisfied with the size of the disk group, click **Next** at the bottom of the page.
- 10 Specify the size of the first virtual disk to be created in the new disk group in the **New virtual disk capacity** box.
- 11 Enter a name for the virtual disk in the **Name** text box.
- 12 Specify the type of files that will be stored on the virtual disk. MD Storage Manager will optimize the virtual disk based on your selection. Your choices include:
 - File system (typical)
 - Database
 - Multimedia
- 13 When you are satisfied with the parameters of the virtual disk, click **Next**.

- 14** To map the new virtual disk to a host now, select **Map now** and assign a logical unit number (LUN) to the virtual disk in the drop-down box.
To map the new virtual disk to a host later, select **Map later**.
- 15** Click **Finish** to create the new disk group and the first virtual disk in the group.

Creating a Virtual Disk From Free Capacity

To manually create a virtual disk within an existing disk group, first decide in which disk group you want to create the new virtual disk. Then use the following procedure to create the new virtual disk:

- 1** Click the **Configure** tab, then click **Create Virtual Disks**.
- 2** Select **Free capacity** on the **Create Virtual Disks – Select Capacity Type** page.
- 3** Click the plus sign (+) at the left of the disk group to display the virtual disks and free capacity in the disk group.
- 4** Click the free capacity for the disk group you want to modify, and then click **Next** at the bottom of the page.
- 5** Specify the size of the virtual disk to be created in the **New virtual disk capacity** box.
- 6** Enter a name for the virtual disk in the **Name** text box.
- 7** Specify the type of files that will be stored on the virtual disk. MD Storage Manager will optimize the virtual disk based on your selection. Your choices include:
 - File system (typical)
 - Database
 - Multimedia
- 8** When you are satisfied with the parameters of the virtual disk, click **Next**.
- 9** To map the new virtual disk to a host now, select **Map now** and assign a logical unit number (LUN) to the virtual disk in the drop-down box.
To map the new virtual disk to a host later, select **Map later**.

After you create virtual disks and map them to hosts, you must register the virtual disks with each host. Registration ensures the host recognizes the virtual disks.

If you plan to create multiple virtual disks, wait until you have created all the virtual disks to register them. Waiting prevents you from having to register virtual disks more than once.

Linux



NOTE: You need super-user (Linux) privileges to run the `hot_add` utility.

The `hot_add` utility is installed with the `host-agent` package and is run from the Linux command line. You cannot run the `hot_add` utility using the MD Storage Manager.

Windows

Windows automatically registers virtual disks.



NOTE: Virtual disks mapped to Windows Server 2008 hosts are marked offline by default. To bring the virtual disks online, use the Disk Management MMC (if you are using a GUI version of Windows Server 2008) or use the `DiskPart` utility (if you are using a Core version of Windows Server 2008).

Hot Spare Drive Protection

Hot spare drives in a storage array provide an additional level of protection in case a physical disk fails. Hot spare drives only take over for failed drives in a RAID level 1, 5, or 10 disk group. Using a hot spare drive can be an advantage because it automatically replaces a failed physical disk that is part of a disk group.

You can see whether hot spare protection is currently set by viewing the **Hot Spare Physical Disks** line in the **Hardware Components** area of the **Summary** tab. You can also see the number of standby and in-use hot spares. A *standby hot spare drive* is a physical disk that has been assigned as a hot spare drive and is available to take over for any failed physical disk. An *in-use hot spare drive* is a physical disk that has been assigned as a hot spare drive and is currently taking over for a failed physical disk.

Automatically Configuring Hot Spares

You can choose to allow MD Storage Manager to automatically configure hot spare drives. With automatic configuration, the controller automatically configures the number and type of hot spare drives that will provide optimal

coverage for the storage array. The number and type of hot spare drives is determined based on the number, type, and capacity of physical disks in the storage array.

To add hot spare drive protection using automatic configuration:

- 1 Click the **Configure** tab, then click **Configure Hot Spares**.
- 2 To automatically assign hot spare drives, click **Configure Hot Spares (Automatic)**.
- 3 Click **Assign**.

To remove hot spare drive protection using automatic configuration:

- 1 Click the **Configure** tab, then click **Configure Hot Spares**.
- 2 Click **Configure Hot Spares (Automatic)**.
- 3 Click **Unassign**.
- 4 Click **OK** in the dialog box.

Manually Configuring Hot Spares

You can choose to manually configure hot spare drives for the drive sets in your storage array. With manual configuration, you assign the type and capacity of hot spare protection for individual drives.

- 1 Click the **Configure** tab, then click **Configure Hot Spares**.
- 2 To manually assign hot spare drives, click **Configure Hot Spares (Manual)**.
- 3 In the **Drive sets** list, click the drive you wish to protect, then click **Assign**.
The **Assign Hot Spares** dialog box appears.
- 4 In the **Assign Hot Spares** dialog box, click the unassigned drive you want to configure into a hot spare physical disk, then click **OK**.



NOTE: When manually configuring a hot spare physical disk, you must use a physical disk type that matches the other physical disks in the disk group. Using a SAS physical disk to replace a SATA II physical disk (or a SATA II physical disk to replace a SAS physical disk) is not supported. Also, the hot spare physical disk must be as large as or larger than the largest physical disk in the disk group.

Host-to-Virtual Disk Mapping

After you create virtual disks, you must map them to the host(s) connected to the array. When you configure host-to-virtual disk mapping, consider these guidelines:

- Each virtual disk in the storage array can be mapped to only one host or host group.
- Host-to-virtual disk mappings are shared between controllers in the storage array.
- A unique LUN must be used by a host group or host to access a virtual disk.
- Not every operating system has the same number of LUNs available.

Creating Host-to-Virtual Disk Mappings

Create host-to-virtual disk mappings by clicking the **Configure** tab, then clicking **Create Host-to-Virtual Disk Mappings**. When you click this link, the Storage Manager displays a series of pages in which you select the hosts and virtual disks to be mapped.

After you complete this configuration, verify the mapping by clicking **Host-to-Virtual Disk Mappings** on the **Summary** tab to ensure the configuration was created correctly.

Modifying and Removing Host-to-Virtual Disk Mapping

You might choose to modify or remove a host-to-virtual disk mapping for several reasons, such as an incorrect mapping or reconfiguration of the storage array. Modifying or removing a host-to-virtual disk mapping applies to both hosts and host groups.



NOTICE: Before you modify or remove a host-to-virtual disk mapping, you must stop any data access (I/O) to the virtual disks to prevent data loss.

- 1 Stop any data access (I/O) to the virtual disks.
- 2 Click the **Modify** tab and then click **Edit Host-to-Virtual Disk Mappings**. MD Storage Manager displays a list of virtual disks and the hosts to which they are mapped.
- 3 Select the virtual disk you wish to modify by clicking its name.

- 4 To map the disk to a different host or host group, click **Change** located to the right of the list.
- 5 To remove the disk mapping to a host or host group, click **Remove** located to the right of the list.

Changing Controller Ownership of the Virtual Disk

If the host has a single data-path to the MD storage array, the virtual disk must be owned by the controller to which the host is connected. You must configure this storage array *before* you start I/O operations and *after* the virtual disk is created.

To assign ownership of the virtual disk to the connected controller:

- 1 Click the **Modify** tab and then select **Change Virtual Disk Ownership/Preferred Path**.
- 2 Select the appropriate virtual disk and click **Change**.

Storage Partitioning

A storage partition is a logical entity consisting of one or more virtual disks that can be accessed by a single host or shared among hosts that are part of a host group. The first time you map a virtual disk to a specific host or host group, a storage partition is created. Subsequent virtual disk mappings to that host or host group do not create another storage partition.

One storage partition is sufficient if:

- Only one attached host accesses all of the virtual disks in the storage array.
- All attached hosts share access to all of the virtual disks in the storage array. When you choose this type of configuration, all of the hosts must have the same operating system and special software (such as clustering software) to manage virtual disk sharing and accessibility.

More than one storage partition is required if:

- Specific hosts must access specific virtual disks in the storage array.
- Hosts with different operating systems are attached to the same storage array. In this case, a storage partition will be created for each host type.



NOTE: If Microsoft Windows is the host server, you must partition and format the virtual disk using the Microsoft Management Console (MMC). Go to **Settings** → **Control Panel** → **Administrative Tools**, or **Control Panel** →

Administrative Tools and select the **Disk Management MMC** option to format and partition the disk. If the host server is running Windows Server 2008 Core version, partition and format the virtual disk using the Microsoft DiskPart utility.

Disk Group and Virtual Disk Expansion

Adding free capacity to a disk group is achieved by adding unconfigured capacity on the array to the disk group. Data is accessible on disk groups, virtual disks, and physical disks throughout the entire modification operation. The additional free capacity can then be used to perform a virtual disk expansion on a standard or snapshot repository virtual disk.

Disk Group Expansion

To add free capacity to a disk group, use the following procedure:

- 1 Click the **Modify** tab, then click **Add Free Capacity (Physical Disks)**.
- 2 On the **Add Free Capacity** page, click the disk group you want to expand.
- 3 Click **Next** beneath the list of disk groups.

The Storage Manager displays information on the disk group you selected.

- 4 In the **Add capacity to volume group** drop-down menu, choose the amount of unconfigured capacity to add to the disk group.
- 5 Click **Finish** to start the process of adding capacity to the disk group.

You can also use the Command Line Interface (CLI) on both Windows and Linux hosts to add free capacity to a disk group. See the *CLI Guide* for more information.

Once the capacity expansion is completed, additional free capacity is available in the disk group for creation of new virtual disks or expansion of existing virtual disks.

Virtual Disk Expansion

Virtual disk expansion is a dynamic modification operation that increases the capacity of standard virtual disks.



NOTE: Snapshot repository virtual disks can be expanded from the CLI or from MD Storage Manager. All other virtual disk types are expandable only from the CLI.

If you receive a warning that the snapshot repository virtual disk is becoming full, you may expand the snapshot repository virtual disk from MD Storage Manager. See "Snapshot Repository Capacity" on page 79 for step-by-step instructions.

Storage Array Media Scan

The media scan is a long-running operation that examines virtual disks to verify that data is accessible. The process finds media errors before normal read and write activity is disrupted and reports errors to the event log.

Errors discovered by the media scan include:

- Unrecovered media error — Data could not be read on the first attempt or on any subsequent attempts. For virtual disks with redundancy protection, data is reconstructed, rewritten to the physical disk, and verified and the error is reported to the event log. For virtual disks without redundancy protection (RAID 0 virtual disks and degraded RAID 1 and RAID 5 virtual disks), the error is not corrected but is reported to the event log.
- Recovered media error — Data could not be read by the physical disk on the first attempt but was successfully read on a subsequent attempt. Data is rewritten to the physical disk and verified and the error is reported to the event log.
- Redundancy mismatches error — The first 10 redundancy mismatches that are found on the virtual disk are reported to the event log.
- Unfixable error — Data could not be read and parity or redundancy information could not be used to regenerate the data. For example, redundancy information cannot be used to reconstruct the data on a degraded virtual disk. The error is reported to the event log.

Changing Media Scan Settings

- 1 Click the Tools tab, then click **Change Media Scan Settings**.
- 2 Select the number of days allowed for the media scan to complete in the **Scan duration (days)** box.



NOTE: Performing the media scan frequently may negatively impact the performance of other operations. Adjust scan duration based on the performance needs of your storage array.

- 3 In the **Select virtual disks to scan** box, click the virtual disk you want to include in the media scan.



NOTE: Press <Ctrl> and click to add more than one virtual disk to the media scan. Click **Select All** to include all virtual disks in the media scan.

- 4 Check the **Scan selected virtual disks** checkbox to enable scanning, then choose either **With consistency check** or **Without consistency check**.

Consistency check enables parity data to be checked during the media scan.

- 5 Click **OK** to accept the updated media scan settings.

Suspending the Media Scan

You cannot perform a media scan while performing another long-running operation on the disk drive such as reconstruction, copy-back, reconfiguration, volume initialization, or immediate availability formatting. If you want to perform another long-running operation, you should suspend the media scan.



NOTE: A background media scan is the lowest priority of the long-running operations.

- 1 Click the **Tools** tab, then click **Change Media Scan Settings**.
- 2 Check the **Suspend media scan** checkbox.
- 3 Click **OK** to suspend media scanning.

Microsoft Services

Virtual Disk Service

The Microsoft Virtual Disk Service (VDS) is supported on your RAID storage array. Microsoft VDS is a set of application programming interfaces (APIs) that provides a single interface for managing disks and other storage hardware, including creating volumes on those disks.

Volume Shadow-Copy Service

The Microsoft Volume Shadow-copy Service (VSS) is a storage management interface for Microsoft Windows operating systems. VSS enables your storage array to interact with third-party applications that use the VSS Application Programming Interface.

Virtual disks that will be used as source virtual disks for VSS snapshots should have names no longer than 16 characters. The VSS hardware provider uses the source virtual disk name as a prefix for the snapshot and repository virtual disk names. The resulting snapshot and repository names will be too long if the source virtual disk name exceeds 16 characters.

For more information on VDS and VSS, see www.microsoft.com.

Premium Feature—Snapshot Virtual Disks



NOTE: If you ordered this feature, you received a Premium Feature Activation card shipped in the same box as your Dell PowerVault MD storage array. Follow the directions on the card to obtain a key file and to enable the feature.

A snapshot virtual disk is a point-in-time image of a virtual disk in a storage array. It is not an actual virtual disk containing data; rather, it is a reference to the data that was contained on a virtual disk at a specific time. A snapshot virtual disk is the logical equivalent of a complete physical copy. However, you can create a snapshot virtual disk much faster than a physical copy, using less disk space.

The virtual disk on which the snapshot is based, called the *source virtual disk*, must be a standard virtual disk in your storage array. Typically, you create a snapshot so that an application, such as a backup application, can access the snapshot and read the data while the source virtual disk remains online and accessible.




NOTE: No I/O requests are permitted on the source virtual disk while the virtual disk snapshot is being created.


A snapshot repository virtual disk containing metadata and copy-on-write data is automatically created when a snapshot virtual disk is created. The only data stored in the snapshot repository virtual disk is that which has changed since the time of the snapshot.

After the snapshot repository virtual disk is created, I/O write requests to the source virtual disk resume. Before a data block on the source virtual disk is modified, however, the contents of the block to be modified are copied to the snapshot repository virtual disk for safekeeping. Because the snapshot repository virtual disk stores copies of the original data in those data blocks, further changes to those data blocks write only to the source virtual disk. The snapshot repository uses less disk space than a full physical copy, because the only data blocks that are stored in the snapshot repository virtual disk are those that have changed since the time of the snapshot.

When you create a snapshot virtual disk, you specify where to create the snapshot repository virtual disk, its capacity, and other parameters. You can disable or delete the snapshot virtual disk when you no longer need it, such as when the backup is complete. If you disable a snapshot virtual disk, you can re-create and reuse it the next time you perform a backup (see "Re-creating Snapshot Virtual Disks" on page 80 for more information). If you delete a snapshot virtual disk, you also delete the associated snapshot repository virtual disk.

 **NOTE:** Deleting a snapshot does not affect data on the source virtual disk.

The information that follows will better prepare users for using the snapshot virtual disk premium feature of the Dell PowerVault systems.

 **NOTE:** The following host preparation sections also apply when using the snapshot feature through the CLI interface.

Creating a Snapshot Virtual Disk Using the Simple Path

You can choose the simple path to create a snapshot virtual disk if the disk group of the source virtual disk has the required amount of free space. A snapshot repository virtual disk requires a minimum 8 MB of free capacity. The destination of a snapshot repository virtual disk is determined based on the free capacity available in the disk group.

If 8 MB of free capacity is not available in the disk group of the source virtual disk, the **Create Snapshot Virtual Disks** feature defaults to the advanced path (see "Creating a Snapshot Virtual Disk Using the Advanced Path" on page 73). In the advanced path option, you can choose to place the snapshot repository virtual disk in another disk group or you can use unconfigured capacity on the storage array to create a new disk group.

About the Simple Path

Using the simple path, you can specify the following parameters for your snapshot virtual disk:

- **Snapshot Virtual Disk Name** — A user-specified name that helps you associate the snapshot virtual disk to its corresponding snapshot repository virtual disk and source virtual disk.

- **Snapshot Repository Virtual Disk Name** — A user-specified name that helps you associate the snapshot repository virtual disk to its corresponding snapshot virtual disk and source virtual disk.

Using the simple path, the following defaults are used for the other parameters of a snapshot virtual disk:

- **Capacity Allocation** — The snapshot repository virtual disk is created using free capacity on the same disk group where the source virtual disk resides.
- **Host-to-Virtual Disk Mapping** — The default setting is **Map now**.
- **Percent Full** — When the snapshot repository virtual disk reaches the specified repository full percentage level, the event is logged in the Main Event Log (MEL). The default snapshot repository full percentage level is 50 percent of the source virtual disk.
- **Snapshot Repository Virtual Disk Full Conditions** — When the snapshot repository virtual disk becomes full, you are given a choice of failing write activity to the source virtual disk or failing the snapshot virtual disk.

Preparing Host Servers to Create the Snapshot Using the Simple Path

- ➡ **NOTICE:** Before using the Snapshot Virtual Disks Premium Feature in a Microsoft® Windows® clustered configuration, you must first map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.

Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the Snapshot enabling process is completed can result in the operating system mis-identifying the snapshot virtual disk. This, in turn, can result in data loss on the source virtual disk or an inaccessible snapshot.

For details on mapping the snapshot virtual disk to the secondary node, refer to the *Dell PowerEdge™ Cluster SE600V Systems Installation and Troubleshooting Guide* on support.dell.com

- 📎 **NOTE:** You can create concurrent snapshots of a source virtual disk on both the source disk group and on another disk group.

Before creating a Snapshot Virtual Disk, note the following:

- The following types of virtual disks are not valid source virtual disks: snapshot repository virtual disks, snapshot virtual disks, target virtual disks that are participating in a virtual disk copy.



NOTE: Virtual Disk Copy is an Advanced (Premium) feature.

- You cannot create a snapshot of a virtual disk that contains unreadable sectors.
- You must satisfy the requirements of your host operating system for creating snapshot virtual disks. Failure to meet the requirements of your host operating system results in an inaccurate point-in-time image of the source virtual disk or the target virtual disk in a virtual disk copy.



NOTICE: Before you create a new point-in-time image of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer®, to make sure all I/O activity has stopped.



NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the Snapshot.


Before creating a snapshot virtual disk, the host server has to be in the proper state. To ensure that the host server is properly prepared to create a snapshot virtual disk, you can either use an application to carry out this task, or you can perform the following steps:

- 1 Stop all I/O activity to the source.
- 2 Using your Windows system, flush the cache to the source. At the host prompt, type


```
SMrepassist -f <filename-identifier>
```

and press <Enter>. See "SMrepassist Utility" on page 102 for more information.

- 3 Remove the drive letter(s) of the source in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot. If this is not done, the snapshot operation will report that it has completed successfully, but the snapshot data will not be updated properly.

 **NOTE:** Verify that the virtual disk has a status of **Optimal** or **Disabled** by clicking the **Summary** tab and then clicking **Disk Groups & Virtual Disks**.


- 4 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.

 **NOTE:** If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

After your host server has been prepared, see "Creating the Snapshot Using the Simple Path" on page 71 to create the snapshot using the simple path.

If you want to use a snapshot regularly, such as for backups, use the **Disable Snapshot** and **Re-create Snapshot** options to reuse the snapshot. Disabling and re-creating snapshots preserves the existing virtual disk-to-host mappings to the snapshot virtual disk.

Creating the Snapshot Using the Simple Path

 **NOTE:** Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the Snapshot.

After first preparing the host server(s) as specified in the preceding procedure, complete the following steps to create a virtual disk snapshot using the simple path:

- 1 In MD Storage Manager, click the **Configure** tab, and then click **Create Snapshot Virtual Disks**.
- 2 The **Additional Instructions** dialog appears; click **Close** in this dialog to continue.
- 3 Click the plus sign (+) to the left of the disk group to expand it, then click the virtual disk from which you want to create a snapshot.
- 4 Click **Next**.

A **No Capacity Exists** warning appears if there is not enough space in the disk group of the source virtual disk to create the snapshot.

- 5 On the **Create Snapshot Virtual Disks – Select Path** screen, select the **Simple path**.



NOTE: A snapshot repository virtual disk requires 8 MB of free space. If the required free space is not available in the disk group of the source virtual disk, the **Create Snapshot Virtual Disks** feature defaults to the advanced path.

- 6 Click **Next**.
- 7 Type a name for the snapshot in the **Snapshot virtual disk name** text box.
- 8 Type a name for the snapshot repository virtual disk in the **Snapshot repository virtual disk name** text box.
- 9 Click **Next**.
- 10 Choose whether to map the virtual disk to a host or host group now or later:

To map now, select **Map now**, select a host or host group by clicking it, and then assign a LUN.

To map later, select **Map later**.

- 11 Click **Finish** to create the snapshot virtual disk and the associated snapshot repository virtual disk.
- 12 After you have created one or more snapshot virtual disks, mount or reassign a drive letter of the source virtual disk.
- 13 If needed, assign host-to-virtual disk mapping between the snapshot virtual disk and the host operating system that accesses it.



NOTE: In some cases, depending on the host type and any virtual disk manager software in use, the software prevents you from mapping the same host to both a source virtual disk and its associated snapshot virtual disk.

- 14 If you are using a Linux-based system, run the `hot_add` utility to register the snapshot virtual disk with the host operating system.



NOTE: The `hot_add` utility is not needed for Windows.

Creating a Snapshot Virtual Disk Using the Advanced Path

About the Advanced Path

Use the advanced path to choose whether to place the snapshot repository virtual disk on free capacity or unconfigured capacity and to change the snapshot repository virtual disk parameters. You can select the advanced path regardless of whether you use free capacity or unconfigured capacity for the snapshot virtual disk.

Using the advanced path, you can specify the following parameters for your snapshot virtual disk:

- **Snapshot Virtual Disk Name** — A user-specified name that helps you associate the snapshot virtual disk to its corresponding snapshot repository virtual disk and source virtual disk.
- **Snapshot Repository Virtual Disk Name** — A user-specified name that helps you associate the snapshot repository virtual disk to its corresponding snapshot virtual disk and source virtual disk.
- **Capacity Allocation** — This parameter allows you to choose where to create the snapshot repository virtual disk. You can allocate capacity by using one of the following methods:
 - Use free capacity on the same disk group where the source virtual disk resides.
 - Use free capacity on another disk group.
 - Use unconfigured capacity and create a new disk group for the snapshot repository virtual disk.

Dell recommends placing the snapshot repository virtual disk within the disk group of the source virtual disk. This ensures that if drives associated with the disk group are moved to another storage array, all the virtual disks associated with the snapshot virtual disk remain in the same group.

- **Percent Full** — When the snapshot repository virtual disk reaches the user-specified repository full percentage level, the event is logged in the Major Event Log (MEL). The default snapshot repository full percentage level is 50% of the source virtual disk.

- **Snapshot Repository Virtual Disk Full Conditions** — You can choose whether to fail writes to the source virtual disk or fail the snapshot virtual disk when the snapshot repository virtual disk becomes full.
- **Host-to-Virtual Disk Mapping** — Choose whether to map the snapshot virtual disk to a host or host group now or to map the snapshot virtual disk later. The default setting is **Map later**.

Preparing Host Servers to Create the Snapshot Using the Advanced Path



NOTICE: Before using the Snapshot Virtual Disks Premium Feature in a Microsoft® Windows® clustered configuration, you must first map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.

Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the Snapshot enabling process is completed can result in the operating system mis-identifying the snapshot virtual disk. This, in turn, can result in data loss on the source virtual disk or an inaccessible snapshot.

For details on mapping the snapshot virtual disk to the secondary node, refer to the *Dell PowerEdge™ Cluster SE600W Systems Installation and Troubleshooting Guide* on support.dell.com

The destination of a snapshot repository virtual disk is determined based on the free capacity available in the disk group. A snapshot repository virtual disk requires a minimum 8 MB of free capacity. You can choose your preferred creation path—simple or advanced—if the disk group of the source virtual disk has the required amount of free space.

If 8 MB of free capacity is not available in the disk group of the source virtual disk, the **Create Snapshot Virtual Disks** feature defaults to the advanced path (see "Creating a Snapshot Virtual Disk Using the Advanced Path" on page 73). In the advanced path option, you can choose to place the snapshot repository virtual disk in another disk group or you can use unconfigured capacity on the storage array to create a new disk group.



NOTE: You can create concurrent snapshots of a source virtual disk on both the source disk group and on another disk group.

Before creating a Snapshot Virtual Disk, note the following:

- The following types of virtual disks are not valid source virtual disks: snapshot repository virtual disks, snapshot virtual disks, target virtual disks that are participating in a virtual disk copy.



NOTE: Virtual Disk Copy is an Advanced (Premium) feature.

- You cannot create a snapshot of a virtual disk that contains unreadable sectors.
- You must satisfy the requirements of your host operating system for creating snapshot virtual disks. Failure to meet the requirements of your host operating system results in an inaccurate point-in-time image of the source virtual disk or the target virtual disk in a virtual disk copy.



NOTICE: Before you create a new point-in-time image of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer®, to make sure all I/O activity has stopped.



NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the Snapshot.

Before creating a snapshot virtual disk, the host server has to be in the proper state. Perform the following steps to prepare your host server:

- 1 Stop all I/O activity to the source.
- 2 Using your Windows system, flush the cache to the source. At the host prompt, type

```
SMrepassist -f <filename-identifier>
```


and press <Enter>. See "SMrepassist Utility" on page 102 for more information.

- 3 Remove the drive letter(s) of the source in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot. If this is not done, the snapshot operation will report that it has completed successfully, but the snapshot data will not be updated properly.



NOTE: Verify that the virtual disk has a status of Optimal or Disabled by clicking the **Summary** tab and then clicking **Disk Groups & Virtual Disks**.


- 4 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.

 **NOTE:** If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

After your host server has been prepared, see "Creating the Snapshot Using the Advanced Path" on page 76 to create the snapshot using the advanced path.

If you want to use a snapshot regularly, such as for backups, use the Disable Snapshot and Re-create Snapshot options to reuse the snapshot. Disabling and re-creating snapshots preserves the existing virtual disk-to-host mappings to the snapshot virtual disk.

Creating the Snapshot Using the Advanced Path


 **NOTE:** Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the Snapshot.

After first preparing the host server(s) as specified in the preceding procedure, complete the following steps to create a virtual disk snapshot using the advanced path:

- 1 In MD Storage Manager, click the **Configure** tab, and then click **Create Snapshot Virtual Disks**.
- 2 The **Additional Instructions** dialog appears; click **Close** in this dialog to continue.
- 3 Click the plus sign (+) to the left of the disk group to expand it, then click the virtual disk from which you want to create a snapshot.
- 4 Click **Next**.

A **No Capacity Exists** warning appears if there is not enough space in the disk group of the source virtual disk to create the snapshot.

- 5 On the **Create Snapshot Virtual Disks – Select Path** screen, select the **Advanced** path.

 **NOTE:** A snapshot repository virtual disk requires 8 MB of free space. If the required free space is not available in the disk group of the source virtual disk, the **Create Snapshot Virtual Disks** feature defaults to the advanced path.

- 6 Click **Next**.
- 7 Type a name for the snapshot in the **Snapshot virtual disk name** text box.
- 8 Type a name for the snapshot repository virtual disk in the **Snapshot repository virtual disk name** text box.
- 9 Click **Next**.
- 10 Choose whether to create the snapshot virtual disk from unconfigured capacity or free capacity.
To create the snapshot virtual disk from unconfigured capacity:
 - a Select **Unconfigured capacity**, then click **Next**.
 - b On the **Create Snapshot Virtual Disks – Specify Capacity** screen, choose a RAID level, then click **Next**.To create the snapshot virtual disk from free capacity:
 - a Select **Free capacity**.
 - b Select a free capacity node, then click **Next**.
- 11 On the **Create Snapshot Virtual Disks – Specify Repository Virtual Disk Capacity** screen, enter how much space you want to allocate for the snapshot repository virtual disk in the **Specify capacity** box, then click **Next**.
- 12 In the **Notify me when the repository disk becomes x% full** box, specify at what percentage MD Storage Manager should consider the snapshot repository virtual disk to be full.
- 13 Specify what should happen if the snapshot repository virtual disk becomes full.
Select **Fail the snapshot virtual disk** to leave the source disk available.
Select **Fail writes to the source virtual disk** to leave the snapshot virtual disk available and stop data from writing to the source virtual disk.
- 14 Choose whether to map the virtual disk to a host or host group now or later.
To map now, select **Map now**, select a host or host group by clicking it, then assign a LUN.

To map later, select **Map later**.

- 15 Click **Finish** to create the snapshot virtual disk and the associated snapshot repository virtual disk.
- 16 After you have created one or more snapshot virtual disks, mount or reassign a drive letter of the source virtual disk.
- 17 If needed, assign host-to-virtual disk mapping between the snapshot virtual disk and the host operating system that accesses it.



NOTE: In some cases, depending on the host type and any virtual disk manager software in use, the software prevents you from mapping the same host to both a source virtual disk and its associated snapshot virtual disk.

- 18 If you are using a Linux-based system, run the `hot_add` utility to register the snapshot virtual disk with the host operating system.



NOTE: The `hot_add` utility is not needed for Windows.

Specifying Snapshot Virtual Disk Names

Choose a name that helps you associate the snapshot virtual disk and snapshot repository virtual disk with its corresponding source virtual disk. The following information is useful when naming virtual disks:

- By default, the snapshot name is shown in the **Snapshot virtual disk name** field as:

`<source-virtual disk-name>-<sequence-number>`

where *sequence-number* is the chronological number of the snapshot relative to the source virtual disk.

The default name for the associated snapshot repository virtual disk that is shown in the **Snapshot repository virtual disk** field is:

`<source-virtual disk-name>-R<sequence-number>`

For example, if you are creating the first snapshot virtual disk for a source virtual disk called **Accounting**, the default snapshot virtual disk is **Accounting-1**, and the associated snapshot repository virtual disk default name is **Accounting-R1**. The default name of the next snapshot virtual disk you create based on Accounting is **Accounting-2**, with the corresponding snapshot repository virtual disk named as **Accounting-R2** by default.

- Whether you use the software-supplied sequence number that (by default) populates the **Snapshot virtual disk name** or the **Snapshot repository virtual disk name** field, the next default name for a snapshot or snapshot repository virtual disk still uses the sequence number determined by the software. For example, if you give the first snapshot of source virtual disk **Accounting** the name **Accounting-8**, and do not use the software-supplied sequence number of 1, the default name for the next snapshot of **Accounting** is still **Accounting-2**.
- The next available sequence number is based on the number of existing snapshots of a source virtual disk. If you delete a snapshot virtual disk, its sequence number becomes available again.
- You must choose a unique name for the snapshot virtual disk and the snapshot repository virtual disks, or an error message is displayed.
- Names are limited to 30 characters. After you reach this limit in either the **Snapshot virtual disk name** or the **Snapshot repository virtual disk name** fields, you can no longer type in the field. If the source virtual disk is 30 characters, the default names for the snapshot and its associated snapshot repository virtual disk use the source virtual disk name truncated enough to add the sequence string. For example, for **Host Software Engineering Group GR-1**, the default snapshot name is **Host Software Engineering GR-1**, and the default repository name would be **Host Software Engineering GR-R1**.

Snapshot Repository Capacity

If you receive a warning that the capacity for the snapshot repository virtual disk is approaching its threshold, you can increase the capacity of a snapshot repository virtual disk by using one of the following methods:

- Use the free capacity available on the disk group of the snapshot repository virtual disk.
- Add unconfigured capacity to the disk group of the snapshot repository virtual disk. Use this option when no free capacity exists on the disk group.

You cannot increase the storage capacity of a snapshot repository virtual disk if the snapshot repository virtual disk has any one of the following conditions:

- The virtual disk has one or more hot spare drives in use.
- The virtual disk has a status other than Optimal.

- Any virtual disk in the disk group is in any state of modification.
- The controller that has ownership of this virtual disk is currently adding capacity to another virtual disk. Each controller can add capacity to only one virtual disk at a time.
- No free capacity exists in the disk group.
- No unconfigured capacity is available to add to the disk group.

To expand the snapshot repository virtual disk from MD Storage Manager:

- 1** Click the **Modify** tab, then click **Modify snapshot virtual disks**.
- 2** Click **Expand Snapshot Repository**.
- 3** Click the snapshot repository virtual disk you want to expand.
- 4** If necessary, you can add free capacity to the volume group by adding an unassigned drive. To add an unassigned drive:
 - a** Click **Add Drives**.
 - b** Select the capacity to add from the drop-down menu.
 - c** Click **Add**.
- 5** Enter the amount by which you want to expand the snapshot repository virtual disk in the **Increase capacity by** field.
- 6** Click **Finish** to expand the capacity of the snapshot repository virtual disk.

Re-creating Snapshot Virtual Disks

Before re-creating a snapshot virtual disk, refer to the following guidelines.

Disabling a Snapshot Virtual Disk

Disable a snapshot virtual disk if one of the following conditions exists:

- You do not need the snapshot now.
- You intend to re-create the snapshot at a later time and want to retain the associated snapshot repository virtual disk so that you do not need to create it again.
- You want to maximize storage array performance by stopping copy-on-write activity to the snapshot repository virtual disk.

The `SMdevices` utility displays the snapshot virtual disk in its output, even after the snapshot virtual disk is disabled.

To disable a snapshot virtual disk:

- 1 Click the **Modify** tab, then click **Modify snapshot virtual disks**.
- 2 Click **Disable Snapshot Virtual Disks**.
- 3 Highlight the snapshot virtual disk to be disabled and click **Disable** beneath the list.
- 4 In the **Confirm Disable Snapshot Virtual Disk** dialog box, type **yes** and then click **OK**.

The snapshot virtual disk is disabled. The associated snapshot repository virtual disk does not change status, but copy-on-write activity to the disabled snapshot virtual disk stops until the snapshot virtual disk is re-created.

Preparing Host Servers to Re-create a Snapshot Virtual Disk



NOTICE: Before you create a new point-in-time image of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk and snapshot virtual disk to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.



NOTE: Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the Snapshot.

Before re-creating a snapshot virtual disk, both the host server and the associated virtual disk you are re-creating have to be in the proper state. Perform the following steps to prepare your host server and virtual disk:

- 1 Stop all I/O activity to the source and snapshot virtual disk (if mounted).
- 2 Using your Windows system, flush the cache to both the source and the snapshot virtual disk (if mounted). At the host prompt, type

```
SMrepassist -f <filename-identifier>
```

and press <Enter>. See "SMrepassist Utility" on page 102 for more information.
- 3 Click the **Summary** tab, then click **Disk Groups & Virtual Disks** to ensure that the snapshot virtual disk is in Optimal or Disabled status.

- 4 Remove the drive letter(s) of the source and (if mounted) snapshot virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot. If this is not done, the snapshot operation will report that it has completed successfully, but the snapshot data will not be updated properly.
- 5 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.



NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Re-creating a Snapshot Virtual Disk

After first preparing the host server(s) as specified in the preceding procedure, re-create a snapshot virtual disk using the following steps.





NOTICE: This action invalidates the current snapshot.


- 1 Click the **Modify** tab, then click **Modify snapshot virtual disks**.
- 2 Click **Re-create Snapshot Virtual Disks**.
- 3 Highlight the snapshot virtual disk to re-create and click **Re-Create** beneath the list.
- 4 In the **Confirm Snapshot Virtual Disk Re-Creation** dialog box, type *yes* and then click **OK**.

Re-creating a snapshot repository virtual disk uses the previously configured snapshot name and parameters.

Premium Feature—Virtual Disk Copy

 **NOTICE:** A virtual disk copy overwrites data on the target virtual disk. Before starting a virtual disk copy, ensure that you no longer need the data or back up the data on the target virtual disk.

 **NOTE:** If you ordered this feature, you received a Premium Feature Activation card that shipped in the same box as your Dell PowerVault MD storage array. Follow the directions on the card to obtain a key file and to enable the feature.

 **NOTE:** The preferred method for creating a virtual disk copy is to copy from a snapshot virtual disk. This allows the original virtual disk used in the snapshot operation to remain fully available for read/write activity while the snapshot is used as the source for the virtual disk copy operation.

When you create a virtual disk copy, you create a copy pair that has a source virtual disk and a target virtual disk on the same storage array.

The *source virtual disk* is the virtual disk that contains the data you want to copy. The source virtual disk accepts the host I/O read activity and stores the data until it is copied to the target virtual disk. The source virtual disk can be a standard virtual disk, a snapshot virtual disk, or the source virtual disk of a snapshot virtual disk. When you start a virtual disk copy, all data is copied to the target virtual disk, and the source virtual disk permissions are set to read-only until the virtual disk copy is complete.

The *target virtual disk* is a virtual disk to which you copy data from the source virtual disk. The target virtual disk can be a standard virtual disk, or the source virtual disk of a failed or disabled snapshot virtual disk.

After the virtual disk copy is complete, the source virtual disk becomes available to host applications for write requests. To prevent error messages, do not attempt to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy is in progress.

Reasons to use virtual disk copy include the following:

- Copying data for improved access — As your storage requirements for a virtual disk change, you can use a virtual disk copy to copy data to a virtual disk in a disk group that uses drives with larger capacity within the same storage array. Copying data for larger access capacity enables you to move data to greater capacity physical disks (for example, 61 GB to 146 GB).
- Restoring snapshot virtual disk data to the source virtual disk — The Virtual Disk Copy feature enables you first to restore the data from a snapshot virtual disk and then to copy the data from the snapshot virtual disk to the original source virtual disk.
- Creating a backup copy — The Virtual Disk Copy feature enables you to create a backup of a virtual disk by copying data from one virtual disk (the source virtual disk) to another virtual disk (the target virtual disk) in the same storage array, minimizing the time that the source virtual disk is unavailable to host write activity. You can then use the target virtual disk as a backup for the source virtual disk, as a resource for system testing, or to copy data to another device, such as a tape drive or other media.



NOTE: Recovering from a backup copy — You can use the Edit Host-to-Virtual Disk Mappings feature to recover data from the backup virtual disk you created in the previous procedure. The Mappings option enables you to unmap the source virtual disk from its host and then to map the backup virtual disk to the same host.

Creating a Virtual Disk Copy for an MSCS Shared Disk

To create a virtual disk copy for a Microsoft Cluster Server (MSCS) shared disk, create a snapshot of the virtual disk, and then use the snapshot virtual disk as the source for the virtual disk copy.



NOTE: An attempt to directly create a virtual disk copy for an MSCS shared disk, rather than using a snapshot virtual disk, will fail with the following error: The operation cannot complete because the selected virtual disk is not a source virtual disk candidate.



NOTE: When creating a snapshot virtual disk, map the snapshot virtual disk to only one node in the cluster. Mapping the snapshot virtual disk to the host group or both nodes in the cluster may cause data corruption by allowing both nodes to concurrently access data.

Virtual Disk Read/Write Permissions

After the virtual disk copy is complete, the target virtual disk automatically becomes read-only to the hosts. The target virtual disk rejects read and write requests while the virtual disk copy operation has a status of Pending or In Progress or if the operation fails before completing the copy. Keep the target virtual disk **Read-Only** enabled if you want to preserve the data on the target virtual disk for reasons such as the following:

- If you are using the target virtual disk for backup purposes.
- If you are using the data on the target virtual disk to copy back to the source virtual disk of a disabled or failed snapshot virtual disk.

If you decide not to preserve the data on the target virtual disk after the virtual disk copy is complete, change the write protection setting for the target virtual disk to **Read/Write**.

To set the target virtual disk read/write permissions, complete the following steps:

- 1 Click the **Modify** tab, and then click **Manage Virtual Disk Copies**.
- 2 Select one or more copy pairs in the table and click **Permissions** to the right of the table.
The **Set Target Virtual Disk Permissions** dialog box appears.
- 3 In the **Set Target Virtual Disk Permissions** dialog box select either **Read-Only** or **Read/Write**.
- 4 Click **OK** in the dialog box.

If you select **Read-Only**, write requests to the target virtual disk will be rejected. If you select **Read/Write**, the host can read and write to the target virtual disk after the virtual disk copy is complete.

Virtual Disk Copy Restrictions

Before you perform any virtual disk copy tasks, understand and adhere to the restrictions listed in this section. The restrictions apply to the source virtual disk, the target virtual disk, and the storage array.

- While a virtual disk copy has a status of In Progress, Pending, or Failed, the source virtual disk is available for read I/O activity only. After the virtual disk copy is complete, read and write I/O activity to the source virtual disk are permitted.
- A virtual disk can be selected as a target virtual disk for only one virtual disk copy at a time.
- A virtual disk copy for any virtual disk cannot be mounted on the same host as the source virtual disk.
- Windows does not allow a drive letter to be assigned to a virtual disk copy.
- A virtual disk with a Failed status cannot be used as a source virtual disk or target virtual disk.
- A virtual disk with a Degraded status cannot be used as a target virtual disk.
- A virtual disk participating in a modification operation cannot be selected as a source virtual disk or target virtual disk. Modification operations include the following:
 - Capacity expansion
 - RAID-level migration
 - Segment sizing
 - Virtual disk expansion
 - Defragmenting a virtual disk



NOTE: The following host preparation sections also apply when using the virtual disk copy feature through the CLI interface.

Creating a Virtual Disk Copy

Use the **Create Virtual Disk Copies** feature on the **Configure** tab to create a full copy of a source virtual disk. This operation overwrites any existing data on the target virtual disk. Once the virtual disk copy has started, all I/O activity to the source virtual disk is read-only. Any attempts to write to the source virtual disk fail until the operation is complete.



NOTE: It is recommended that you create a virtual disk copy from a snapshot virtual disk rather than from the original virtual disk. This allows the original virtual disk to remain in full use while the snapshot of this virtual disk is used as the source for the virtual disk copy operation.

Preparing Host Servers to Create a Virtual Disk Copy



NOTICE: Before you create a new copy of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk (and, if applicable, the target disk) to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.



NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the virtual disk copy.

Before creating a virtual disk copy, both the host server and the associated virtual disk you are copying have to be in the proper state. Perform the following steps to prepare your host server and virtual disk:

- 1 Stop all I/O activity to the source and target virtual disk.
- 2 Using your Windows system, flush the cache to both the source and the target virtual disk (if mounted). At the host prompt, type
`SMrepassist -f <filename-identifier>`
and press <Enter>. See "SMrepassist Utility" on page 102 for more information.
- 3 Click the **Summary** tab, then click **Disk Groups & Virtual Disks** to ensure that the virtual disk is in Optimal or Disabled status.

- 4 Remove the drive letter(s) of the source and (if mounted) virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the virtual disk. If this is not done, the copy operation will report that it has completed successfully, but the copied data will not be updated properly.
- 5 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable virtual disk copies.



NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Copying the Virtual Disk

After first preparing the host server(s) as specified in the preceding procedure, complete the following steps to make a virtual disk copy:

- 1 Click the **Configure** tab, then click **Create Virtual Disk Copies**.
- 2 On the **Select Source Virtual Disk** page, select the virtual disk to copy (source virtual disk), and click **Next**.

The **Select Target Virtual Disk** page appears.



NOTE: If the virtual disk you select is not valid, an information dialog box appears explaining the types of virtual disks you can use as the source for a virtual disk copy. Click **OK** to close this dialog box and select a different source virtual disk.

- 3 Choose the target virtual disk:
 - a To use an existing virtual disk as the target, select **Use an existing virtual disk** and highlight the virtual disk of your choice in the list.
- 4 Click **Next** at the bottom of the page.



NOTE: If you select a target virtual disk with a capacity similar to the source virtual disk, you reduce the risk of having unusable space on the target virtual disk after the virtual disk copy is completed.

- b To create a new virtual disk for the target, select **Create a new virtual disk**. Type a name for this new target virtual disk in the text box.

The **Create virtual disk copies—Set Copy Priority** dialog box appears.

- 5 Set the copy priority for the virtual disk copy and click **Next**.

The source virtual disk, the target virtual disk, and the copy priority setting that you selected appear on the **Create virtual disk copies—Confirm Copy Settings** dialog. The higher priorities allocate more resources to the virtual disk copy at the expense of the storage array's performance. For more information, see "Setting Copy Priority" on page 89.

Storage Array Performance During Virtual Disk Copy

The following factors contribute to the overall performance of the storage array:

- I/O activity
- Virtual disk RAID level
- Virtual disk configuration — Number of drives in the virtual disk groups
- Virtual disk type — Snapshot virtual disks might take more time to copy than standard virtual disks

During a virtual disk copy, resources for the storage array are diverted from processing I/O activity to completing a virtual disk copy. This affects the overall performance of the storage array. When you create a new virtual disk copy, you define the copy priority to determine how much controller processing time is diverted from I/O activity to a virtual disk copy operation.

Setting Copy Priority

The **Copy Priority** setting defines how much of the storage array's resources are used to complete a virtual disk copy, rather than to fulfill I/O requests. Changing the copy priorities sets the rate at which a virtual disk copy is completed.

Five copy priority rates are available: lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized and the virtual disk copy takes longer. At the highest priority rate, the virtual disk copy is prioritized, and I/O activity for the storage array is slower.

You can change the copy priority for a virtual disk copy in the following circumstances:

- Before the virtual disk copy begins
See "Creating a Virtual Disk Copy" on page 87.
- While the virtual disk copy is in progress
Click the **Modify** tab, then click **Manage Virtual Disk Copies**. Select an active copy operation, then click **Priority** to the right of the list of virtual disk copies.
- When re-creating a virtual disk copy
Click the **Modify** tab, then click **Manage Virtual Disk Copies**. Select a completed copy operation, then click **Re-copy** to the right of the list of virtual disk copies.

Stopping a Virtual Disk Copy

You can stop a virtual disk copy operation that has an In Progress status, a Pending status, or a Failed status. Stopping a virtual disk copy that has a Failed status clears the Needs Attention status displayed for the storage array.

When you stop a virtual disk copy, all mapped hosts have write access to the source virtual disk. If data is written to the source virtual disk, the data on the target virtual disk no longer matches the data on the source virtual disk.

To stop a virtual disk copy, complete the following steps:


- 1 Click the **Modify** tab, and then click **Manage virtual disk copies**.
- 2 Select the copy operation you wish to stop by clicking it and click **Stop**.
You can only select one copy operation at a time to be stopped.
- 3 Click **Yes** to stop the virtual disk copy.


Recopying a Virtual Disk

You can recopy a virtual disk when you have stopped a virtual disk copy and you want to start it again or when a virtual disk copy has failed.

The Recopy option overwrites existing data on the target virtual disk and makes the target virtual disk read-only to hosts. This option fails all snapshot virtual disks associated with the target virtual disk, if any exist.

Preparing Host Servers to Recopy a Virtual Disk

 **NOTICE:** Before you create a new copy of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk (and, if applicable, the target disk) to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.


 **NOTE:** Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux will help to guarantee a stable copy of the drive for the virtual disk copy.

Before creating a new virtual disk copy for an existing copy pair, both the host server and the associated virtual disk you are recopying have to be in the proper state. Perform the following steps to prepare your host server and virtual disk:

- 1 Stop all I/O activity to the source and target virtual disk.
- 2 Using your Windows system, flush the cache to both the source and the target virtual disk (if mounted). At the host prompt, type

```
SMrepassist -f <filename-identifier>
```

and press <Enter>. See "SMrepassist Utility" on page 102 for more information.
- 3 Click the **Summary** tab, then click **Disk Groups & Virtual Disks** to ensure that the virtual disk is in Optimal or Disabled status.
- 4 Remove the drive letter(s) of the source and (if mounted) virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the virtual disk. If this is not done, the copy operation will report that it has completed successfully, but the copied data will not be updated properly.
- 5 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable virtual disk copies.

 **NOTE:** If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Recopying the Virtual Disk


After first preparing the host server(s) as specified in the preceding procedure, complete the following steps to create a new virtual disk copy for an existing copy pair:

- 1 Click the **Modify** tab, and then click **Manage virtual disk copies**.
You can only select one copy operation at a time to be recopied.
- 2 Select the copy operation in the list displayed by the **Manage Virtual Disk Copies** page, and then click **Recopy** at the right of the list.
- 3 The **Recopy** dialog box appears. Set the copy priority.
- 4 Type **yes**, and click **OK**.
- 5 If you approve of the parameters, type **yes** in the text box and click **Finish** to confirm the copy settings and start the virtual disk copy.
The **Copy Started** page appears, verifying that the virtual disk copy has started. This dialog also enables you to exit the **Create virtual disk copies** feature or create another new virtual disk copy.
- 6 Choose one of the following options, based on whether you want to create another virtual disk copy or modify the one you just created:
 - **Yes** — Create a new virtual disk copy.
 - **No** — Exit the Create virtual disk copies dialog.
 - **Manage Virtual Disk Copies** — Recopy, stop the copy process, set permissions or priority, or remove virtual disk copies.

You can view the progress of a virtual disk copy in the **Manage virtual disk copies** page. For each copy operation in progress, the list displays a sliding scale in the Status field showing the percentage of the operation that is complete.

Once the virtual disk copy is complete, perform the following actions:

- 1 In Linux, if you created the target virtual disk with unconfigured capacity, run the `hot_add` utility.
- 2 If you created the target virtual disk with unconfigured capacity, you must map the virtual disk to a host in order to use it. See "Host-to-Virtual Disk Mapping" on page 60 for more information.

- 3 You must register the target virtual disk with the operating system before you can use the new virtual disk. Perform the following steps:
 - a Enable write permission on the target virtual disk by either removing the virtual disk copy pair (see "Removing Copy Pairs" on page 93) or explicitly setting write permission.
 - b In Windows, assign a drive letter to the virtual disk.
 **NOTE:** Following a disk copy, if the properties of a Windows-based volume indicate a RAW file system (one that has not been formatted) you must reboot the system so that Windows can recognize the correct virtual target disk.
 - c In Linux, mount the virtual disk.
- 4 Enable I/O activity to the source virtual disk and the target virtual disk.

Removing Copy Pairs

Removing copy pairs permanently removes any virtual disk copy-related information for the source virtual disk and target virtual disk in the **Virtual Disk Properties** and the **Storage Array Profile** dialogs.

After you remove the virtual disk copy, you can select the target virtual disk as a source virtual disk or a target virtual disk for a new virtual disk copy. Removing a virtual disk copy also permanently removes the Read-Only attribute for the target virtual disk.

Removing copy pairs does not delete the data on the source virtual disk or target virtual disk. This merely breaks the copy relationship between the two virtual disks.

When you remove a virtual disk copy from the storage array, the target write attribute for the target virtual disk is also removed. If the virtual disk copy is in In Progress status, you must stop the virtual disk copy before you can remove the copy pair.

To remove a copy pair, perform the following steps:

- 1 Click the **Modify** tab, and then click **Manage virtual disk copies**.
- 2 Select one or more copy pairs in the table, and click **Remove**.
The **Remove Copy Pairs** dialog appears.
- 3 Click **Yes** to remove the copy pair.

Firmware Downloads

You can download the following types of firmware images with MD Storage Manager:

- RAID controller module firmware that manages the storage array controllers
- RAID nonvolatile static random access memory (NVS RAM) images that specify the default settings for the storage array controllers
- Physical disk firmware that controls the operation of the disks in the storage array
- Enclosure Management Modules (EMMs) firmware that manages data transfer between the drives and a RAID enclosure

To download firmware for any of these components, click the **Support** tab and then click **Download firmware**. MD Storage Manager displays links to specify the components for which firmware is available.

Before downloading any firmware, verify that the storage arrays are in Optimal status. When you download new firmware, MD Storage Manager checks the operating status of the storage array controllers. If any controllers are not in the Optimal status, an error message appears, and you can stop or continue the download. Before continuing, correct any non-Optimal conditions.



NOTE: Virtual disks that do not have all their member drives at controller startup are reported as Optimal. The firmware reports this to prevent disks failing as a result of disconnected enclosures.

Downloading RAID Controller and NVSRAM Packages

The following sections describe the downloading process for RAID Controller and NVSRAM firmware.



NOTE: Due to a limitation with Linux, firmware updates to the RAID controller module must be performed using out-of-band management only. Failure to do so may result in the host server becoming unresponsive, and it may require a reboot.

Downloading Both RAID Controller and NVSRAM Firmware



NOTE: I/O to the array can continue while you are upgrading RAID controller and NVSRAM firmware.



NOTE: The RAID enclosure must contain at least two disk drives in order to update the firmware on the controller.

Use the following procedure to download RAID controller and NVSRAM firmware in a single operation:

- 1 Click the **Support** tab, then click **Download firmware**.
- 2 From the **Download firmware** display, click **Download RAID Controller Module Firmware**.

A dialog box lists the current controller firmware and NVSRAM versions in use.

- 3 Click **Select File** to browse to the file that you want to download. By default, only firmware images that are compatible with the current storage array configuration appear.
- 4 Click the file in the **File Selection** area and then click **OK**.
- 5 If the file you selected is not valid or is incompatible with the current storage array configuration, an error message appears. Click **OK** to close the error message, and select a compatible file.



NOTE: If you wish to only download firmware for the RAID controller, skip to step 10 in this procedure.

- 6 Click the check box next to **Transfer NVSRAM file with RAID controller module firmware**.
- 7 Click **Select File** to browse to the file that you want to download. By default, only firmware images that are compatible with the current storage array configuration appear.
- 8 Click the file in the **File Selection** area and then click **OK**.
- 9 If the file you selected is not valid or is incompatible with the current storage array configuration, an error message appears. Click **OK** to close it, and select a compatible file.
- 10 Click **Transfer...**

- 11 A **Confirm Download** dialog box appears listing the current versions and the versions you selected of the RAID controller and NVSRAM firmware. To complete the download, click **Yes**.

Downloading Only NVSRAM Firmware

Use the following procedure to download NVSRAM firmware:

- 1 Click the **Support** tab, then click **Download firmware**.
- 2 From the **Download firmware** display, click **Download RAID Controller Module NVSRAM**.
A dialog box lists the current controller firmware and NVSRAM versions in use.
- 3 Click **Select File** to browse to the file that you want to download. By default, only firmware images that are compatible with the current storage array configuration appear.
- 4 Click the file in the **File Selection** area and then click **OK**.
- 5 If the file you selected is not valid or is incompatible with the current storage array configuration, an error message appears. Click **OK** to close it, and select a compatible file.
- 6 Click **Transfer...**
- 7 A **Confirm Download** dialog box appears listing the current versions and the versions you selected of the RAID controller and NVSRAM firmware. To complete the download, click **Yes**.

Downloading Non-redundant MSCS NVSRAM Firmware



NOTE: For non-redundant MSCS Cluster configurations, download an updated NVSRAM to avoid Virtual Disk Not on Preferred Path conditions.

Use the following procedure to download non-redundant NVSRAM firmware:

- 1 Click the **Support** tab, then click **Download firmware**.
- 2 From the **Download firmware** display, click **Download RAID Controller Module NVSRAM**.
A dialog box lists the current controller firmware and NVSRAM versions in use.

- 3 Click **Select File** to browse to the file that you want to download. By default, only firmware images that are compatible with the current storage array configuration appear.
- 4 Click the **non-redundant-MSCS NVSRAM** file in the **File Selection** area and then click **OK**.
- 5 If the file you selected is not valid or is incompatible with the current storage array configuration, an error message appears. Click **OK** to close it, and select a compatible file.
- 6 Click **Transfer...**
- 7 A **Confirm Download** dialog box appears listing the current versions and the versions you selected of the RAID controller and NVSRAM firmware. To complete the download, click **Yes**.

Downloading Physical Disk Firmware

Use the following procedure to download physical disk firmware:



NOTE: Dell recommends stopping all I/O to the array when downloading physical disk firmware.





NOTE: Due to a limitation with Linux, physical disk firmware updates must be performed using out-of-band management only. Failure to do so may result in the host server becoming unresponsive, and it may require a reboot.

- 1 Click the **Support** tab, then click **Download firmware**.
- 2 From the **Download firmware** display, click **Download Physical Disk Firmware**.
A dialog box lists the current physical disk firmware version in use.
- 3 Click **Add** to browse to the file that you want to download. By default, only firmware images that are compatible with physical disks in the storage array appear.
- 4 Click the file in the **File Selection** area and then click **OK**.
- 5 If the file you selected is not valid or is incompatible with the physical disks in the storage array, an error message appears. Click **OK** to close it, and select a compatible file.
- 6 Click **Transfer...**

- 7 A **Confirm Download** dialog box appears listing the current versions and the versions you selected of physical disk firmware. To complete the download, click **Yes**.

Downloading EMM Firmware

 **NOTICE:** Do not make any configuration changes to the storage array while you are downloading the EMM firmware. Doing so could cause the firmware download to fail, damage the storage array, or cause loss of data accessibility.

 **NOTE:** Due to a limitation with Linux, EMM firmware updates must be performed using out-of-band management only. Failure to do so may result in the host server becoming unresponsive, and it may require a reboot.

- 1 Click the **Support** tab, then click **Download firmware**.
- 2 From the **Download firmware** display, click **Download Environmental (EMM) Card Firmware**.
A list of expansion enclosures appears with the corresponding version of the current EMM firmware file.
- 3 Select where to download the EMM firmware by clicking an individual expansion enclosure to highlight it or by clicking the **Select All** checkbox to highlight all the expansion trays.
- 4 Click **Select File** to locate the directory in which the EMM firmware file to download resides. Select the file to download by double-clicking the file, and then click **Start** to start the download.

The **Start** button is disabled until you select a firmware file. If you click **Stop** while a firmware download is in progress, the download completes before the operation stops. When the status field for the remaining expansion enclosures changes to **Canceled**, restart the firmware upgrade process.

Troubleshooting Problems

The following sections provide information to assist you in resolving problems that may occur with your MD Storage Array.

Recovery Guru

The Recovery Guru is a component of MD Storage Manager that diagnoses critical events on the storage array and recommends step-by-step recovery procedures for problem resolution. You can access the Recovery Guru by clicking **Storage Array Needs Attention** on the **Summary** page or by clicking **Recover from failure** on the **Support** page.

You can detect a problem using the following indicators:

- Non-Optimal status icons
- Alert notification messages that are sent to the appropriate destinations
- Hardware indicator lights

The status icons return to Optimal status as problems are resolved.

Storage Array Profile

The Storage Array Profile provides an overview of your configuration, including firmware versions and the current status of all devices on the storage array.

You can access the storage array profile by clicking **View storage array profile** from either the **Summary** or **Support** pages.

Device Health Conditions

The storage array establishes communication with each managed device and determines the current device status. Before you configure or troubleshoot your device, always make sure the enclosures in the storage array are in Optimal status.

A storage array is always in one of six possible health status conditions, which you can identify by the status icon.

- Optimal status — Every component in the managed device is in the desired working condition.
- Needs Attention status — A problem exists with the managed device that requires intervention. If the storage array has a Needs Attention status, contact Technical Assistance for resolution.
- Fixing status — A Needs Attention condition has been corrected and the managed device is currently moving into Optimal status.
- Unresponsive status — The storage management station cannot communicate with the device or with one or both controllers in the storage array.
- Contacting Device status — MD Storage Manager is currently establishing contact with the device.
- Needs Upgrade status — The storage array is running a level of firmware that is no longer supported by MD Storage Manager.



NOTE: For every non-Optimal status condition listed, use the Recovery Guru to detect and troubleshoot the problem.



NOTE: Wait at least five minutes for the storage array to return to an Optimal status following a recovery procedure.

SMrepassist Utility

SMrepassist (replication assistance) is a host-based utility for Windows platforms. This utility is installed with the MD Storage Manager software. Use this utility before and after you create a virtual disk copy on a Windows operating system to ensure that all the memory-resident data for file systems on the target virtual disk is flushed and that the driver recognizes signatures and file system partitions. You can also use this utility to resolve duplicate signature problems for snapshot virtual disks.

From a MS-DOS[®] window on a host running Windows, navigate to `C:\Program Files\Dell\MD Storage Manager\util` and run the following command:

```
SMrepassist -f <filesystem-identifier>
```

where `-f` flushes all the memory-resident data for the file system indicated by `<filesystem-identifier>`, and `<filesystem-identifier>` specifies a unique file system in the following syntax:

```
drive-letter: <mount-point-path>
```

The file system identifier might consist of only a drive letter, as in the following example:

```
SMrepassist -f E:
```



NOTE: In Windows, the mount point path is a drive letter.

An error message appears in the command line when the utility cannot distinguish between the following:

- Source virtual disk and snapshot virtual disk (for example, if the snapshot virtual disk has been removed)
- Standard virtual disk and virtual disk copy (for example, if the virtual disk copy has been removed)

Support Information Package

MD Storage Manager provides a feature that enables you to save all storage array data, such as profile and event log information, to a file that you can send if you seek technical assistance for problem resolution. To generate this support information bundle:

- 1 Click the **Support** tab, then click **Gather Support Information**.
- 2 Click **Browse**.

The **Collect All Support Data** dialog box appears.

- 3 In the **Save in** drop-down menu, select the location at which you want to save the support data bundle. In the **File name** text box, type a name for the bundle.
- 4 Click **Save** to close the **Collect All Support Data** dialog box.
- 5 Click **Start**.

The support information bundle is saved to the location of your choice.

Unidentified Devices

An unidentified node or device occurs when the MD Storage Manager cannot access a new storage array. Causes for this error include network connection problems, the storage array is turned off, or the storage array does not exist.



NOTE: Before beginning any recovery procedure, make sure that the host-agent software is installed and running. If you started the host before the host was connected to the storage array, the host-agent software will not be able to find the storage array. If so, make sure that the connections are tight, and restart the host-agent software.

- If a storage array is managed by using both out-of-band management and in-band management using the same host, a management network connection problem might prevent direct communication with the storage array. However, you might still be able to manage the storage array over the in-band connections. The opposite situation can also occur.
- If a storage array is managed through more than one host, it is possible that the storage array might become unresponsive to communication over the connections given by one host. However, you might still be able to manage the storage array over the connections provided by another host.

Recovering from an Unidentified Storage Array

Use the following procedure to recover from an unidentified storage array.

- 1 Make sure that the network connection to the storage management station is operating.
- 2 Make sure that the controllers are installed and that the power is turned on to the storage array. Correct any existing problems before continuing.
- 3 If you have an in-band storage array, use the following procedure. Click **Refresh** after each step to check the results:
 - a Make sure that the host-agent software is installed and running. If you started the host before the host was connected to the controllers in the storage array, the host-agent software will not be able to find the controllers. If so, make sure that the connections are tight, and restart the host-agent software.

- b** Make sure that the network can access the host by using the `ping` command in the following syntax:

```
ping <host-name-or-IP-address-of-the-host>.
```

If the network can access the host, continue to step c. If the network cannot access the host, skip to step d.

- c** Remove the host with the unresponsive status from the MD Storage Manager, and add that host again.

If the host returns to optimal status, you have completed this procedure.

- d** Make sure that the power to the host is turned on and that the host is operational.

- e** If applicable, make sure that the host bus adapters have been installed in the host.

- f** Examine all external cables and switches or hubs to make sure that you cannot see any damage and that they are tightly connected.

- g** If you have recently replaced or added the controller, restart the host-agent software so that the new controller is found.

If a problem exists, make the appropriate modifications to the host.

- 4** If you have an out-of-band storage array, use the following procedure. Click **Refresh** after each step to make sure of the results:

- a** Make sure that the network can access the controllers by using the `ping` command. Use the following syntax:

```
ping <controller-IP-address>.
```

If the network can access the controllers, continue to step b. If the network cannot access the controllers, skip to step c.

- b** Remove the storage array with the unresponsive status from the MD Storage Manager, and add that storage array again.

If the storage array returns to optimal status, you have completed this procedure.

- c** Examine the ethernet cables to make sure that you cannot see any damage and that they are tightly connected.

- d Make sure that the applicable network configuration tasks have been done (for example, the IP addresses have been assigned to each controller).
- 5 Make sure that the controller firmware is compatible with the MD Storage Manager on your management station. If the controller firmware was upgraded, the MD Storage Manager might not have access to the storage array. A new version of MD Storage Manager might be needed to manage the storage array with the new version of the controller firmware.

If this problem exists, see the Dell support website at support.dell.com.

- 6 Look to see if there is too much network traffic to one or more controllers. This problem corrects itself because the MD Storage Manager tries to re-establish communication with the controllers in the storage array at regular times. If the storage array was unresponsive and a subsequent try to connect to the storage array succeeds, the storage array becomes responsive.
- 7 For an out-of-band storage array, look to see if management operations are taking place on the storage array from other storage management stations. The type of management operations being done and the number of management sessions taking place together establish the number of TCP/IP connections made to a controller. When the maximum number of TCP/IP connections have been made, the controller stops responding. This problem corrects itself because after some TCP/IP connections complete, the controller then becomes responsive to other connection tries.
- 8 If the storage array is still unresponsive, problems might exist with the controllers.

If these problems exist, see the Dell support website at support.dell.com.

Enclosure Hardware Replacement, Maintenance, and Configuration Considerations

Removing and Inserting Enclosure Management Modules on Attached Expansion Enclosures

The following procedures describe how to safely remove and insert an enclosure management module (EMM) from an expansion enclosure attached to the MD3000/MD3000i.

➡ NOTICE: Failure to follow these guidelines may result in a physical disk failing during removal and/or inadvertent removal of its redundant data path.

Removing an EMM from the Expansion Enclosure

- 1 Check the Recovery Guru to confirm that there is no loss of physical disk path redundancy.
 - If there is no loss of redundancy, check the channel. If it matches the EMM you are removing, you can safely remove it now.
 - If redundancy is lost, run the following command. Note that `channel [1]` is the degraded channel.

```
"set physicalDiskChannel channel [1]  
status = optimal;"
```

- 2 Verify that the path to the channel is restored. The Recovery Guru may take several minutes to update.

Inserting an EMM into an Expansion Enclosure

- 1 Make sure to always insert the EMM without SAS cables attached.
After the EMM is inserted, wait at least 30 seconds before attaching the SAS cables.

- 2 After attaching the SAS cables, wait at least three minutes for the EMM to reach optimal state.



NOTE: You may experience a transitional failure on multiple EMMs in the MEL log while the system updates.

Removing and Inserting Physical Disks

Refer to the following guidelines to ensure that physical disks are safely removed from and inserted into the MD3000/MD3000i RAID storage array.

- Wait at least 60 seconds between removing a drive and inserting a replacement.
- When pulling a drive from a storage array to move it to a different slot, wait 60 seconds before inserting the drive into the new slot.
- Wait at least 60 seconds between the removal of drives from a storage array.
- Wait at least 60 seconds between the insertion of drives into a storage array.

In a large configuration, storage management software may take up to 10 seconds to detect hardware changes.

MD3000 Maintenance Considerations

For Linux kernels, stop and then restart the SMagent after performing one of the following maintenance tasks:

- Move a controller offline or replace a controller.
- Remove SAS cables from or attach SAS cables to host servers running Red Hat® Enterprise Linux® (version 4), SUSE® Linux Enterprise Server 9, or SUSE Linux Enterprise Server 10 operating systems.



NOTE: If a Resolve Topology Conflicts message appears in the Status portlet of the Summary page after restarting the SMagent, it may be necessary to restart the host server to clear this message. Do not select the **Resolve** selection in the Topology Conflict wizard. The MD3000 will continue to service I/O requests, but the array could become partially managed if only in-band management is used. Stop all I/O operations on any host servers that were involved in the maintenance operation and restart the system.

MD3000 Cluster Configuration Guidelines for Standalone Host Servers

If one of the standalone host servers you are planning to configure into a cluster environment is running MD Storage Manager and has a virtual disk mapped to the array, use that host to create the Host Group and quorum virtual disk mapping.



NOTE: Failure to follow this mapping protocol can cause the host server to lose communication with the array if the server is using only in-band management and only one of the servers has a virtual disk mapped to the array. If communication is lost, restore in-band management to the other host server and complete the cluster setup using the following procedure.

- 1 Using MD Storage Manager from either the host server with restored in-band management or from a management station, select **Configure** → **Create Host-to-Virtual Disk Mappings**.
- 2 Select the host server name that does not have a virtual disk mapped.
- 3 Click **Access** to access the virtual disk.
- 4 Assign a logical unit number (LUN) of 31.
- 5 Click **Finish**.

Index

A

- alert
 - e-mail, 27
 - SNMP, 27

B

- backup, 84

C

- Configure tab, 12
- copy pairs
 - removing, 93
- copy priority, 89

D

- disk
 - expansion, 62
- disk group, 12, 53-55, 80
- documentation, 14
- download firmware, 13

E

- e-mail alert, 27
- event monitor, 15

F

- free capacity, 53

H

- HBA port, 43
- host, 15
 - access, 12
 - configuration, 43-59
 - group, 47
 - name, 43
 - type, 43
- host group, 15, 43
- host-to-virtual disk mapping, 12, 60, 69, 74
- hot spare drive, 58, 79
 - automatic configuration, 59
 - in-use, 53
 - manual configuration, 59
 - standby, 53
- hot_add utility, 58

I

iSCSI

- advanced host port settings, 35
- changing target discovery, 34
- changing target identification, 34
- configuring host ports, 35
- edit, remove, or rename host topology, 38
- iSCSI tab, 13
- viewing or ending a session, 36
- viewing or setting statistics, 37

L

logical unit number, 15

M

Modify tab, 12

N

NVSRAM file, 13

P

- password, 22
- physical disk, 15

R

RAID level, 54-55, 89

S

- safety information, 14
- SMrepassist utility, 102
- snapshot repository virtual disk, 53, 67
 - capacity, 79
- snapshot virtual disk, 12, 53, 67, 84
 - advanced path, 73
- SNMP alert, 15
- source virtual disk, 67, 79, 83, 87
- status, 79
- storage array
 - adding, 17-18
 - adding with automatic discovery, 17
 - managing, 13
 - naming, 18
 - removing, 18
- storage partition, 61
- storage partitioning, 53
- Summary tab, 12, 58
- Support tab, 13
- support.dell.com, 14

T

- target virtual disk, 83, 90
- Tools tab, 13

U

unconfigured capacity, 53

utilities

- hot_add, 58

V

virtual disk, 53-54

- access, 15

- deleting in Linux, 55

- name, 78-79

- recopy, 90

- registering, 57

- source, 67, 79

virtual disk copy, 12, 87, 90

- examples, 84

- restrictions, 86

- stop, 90

W

warranty, 14

